arXiv:1402.5391v1 [math.PR] 21 Feb 2014

Exact Simulation for Assemble-To-Order Systems

Ana Bušić* Emilie Coupechoux †

February 24, 2014

Abstract

We develop exact simulation (also known as perfect sampling) algorithms for a family of assemble-to-order systems. Due to the finite capacity, and coupling in demands and replenishments, known solving techniques are inefficient for larger problem instances. We first consider the case with individual replenishments of items, and derive an event based representation of the Markov chain that allows applying existing exact simulation techniques, using the monotonicity properties or bounding chains. In the case of joint replenishments, the state space becomes intractable for the existing methods. We propose new exact simulation algorithms, based on aggregation and bounding chains, that allow a significant reduction of the state space of the Markov chain. We also discuss the coupling times of considered models and provide sufficient conditions for linear (in the single server replenishment case) or quadratic (many server case) complexity of our algorithms in terms of the total capacity in the system.

Keywords: exact simulation; Markov chains; aggregation; assemble-to-order systems

1 Introduction

Exact simulation (also called perfect sampling) algorithms draw unbiased samples from a target distribution. This distribution is in general not known, or cannot be efficiently computed due to:

- the computational complexity of the normalizing constant of an otherwise known distribution (e.g. in statistical physics applications, approximation algorithms of $\sharp P$-complete problems, or analysis of queueing systems with product form solution), or
- the target distribution that is a limiting distribution of a Markov chain that cannot be efficiently solved by analytic methods (e.g. queueing systems not having a product form solution).

We will focus mainly on the latter case and develop exact simulation algorithms for Assemble-To-Order (ATO) systems.

[Propp and Wilson, 1996] used a *coupling from the past* scheme to derive an exact simulation algorithm - called PSA (Perfect Sampling Algorithm) in the following - providing unbiased samples from the stationary distribution of an ergodic Markov chain with a finite state space. Many variants of their algorithm have been developed since in various contexts. We give in Section 2 a brief overview of PSA and mention some works directly linked to the present article (for more information, see the annotated bibliography by [Wilson, 1998]).

PSA provides unbiased samples from the stationary distribution of an ergodic Markov chain in finite expected time. This is its main advantage over Markov Chain Monte Carlo (MCMC) simulation methods that construct one trajectory of the chain and stop after some long enough burn in period, when the distribution of the current state is estimated to be close enough to

*INRIA and Computer Science Department of École Normale Supérieure (DI ENS), Paris, France. E-mail: ana.busic@inria.fr.

†Laboratoire I3S, Université Nice Sophia Antipolis, Paris, France. E-mail: coupecho@i3s.unice.fr. This paper is part of the author's PhD thesis done at INRIA-ENS.

the stationary distribution (see [Asmussen and Glynn, 2007] for an overview of MCMC methods). This burn in period depends on the mixing time of the chain, that is in general difficult to compute. The stopping criterion for MCMC is thus often based on pessimistic bounds, that can be much larger than the coalescence time of the coupling from the past scheme used by PSA (note that the coalescence time provides an upper bound for the mixing time, see for instance [Levin et al., 2009]). Moreover, PSA detects the exact coalescence time; there is no explicit need for its estimation for the exactness of the algorithm. This represents a clear advantage over MCMC methods using stopping criteria based on the mixing time, when the latter is not known.

The efficiency of PSA *a priori* depends on the size of the state space. When the system is monotone, one can easily overcome this issue (as mentioned in [Propp and Wilson, 1996] and explained in Section 2). Developing efficient exact simulation algorithms for non-monotone Markov chains with a very large state space remains challenging: Bounding chains can be used to detect coalescence, as in [Kendall and Møller, 2000], [Huber, 2004] or [Bušić et al., 2012a].

[Kendall and Møller, 2000] describe the general idea and provide the construction of bounding chains for spatial birth-and-death processes. [Huber, 2004] gives a general approach for Markov chains with local interactions, including Gibbs or Metropolis-Hastings samplers. Envelope Perfect Sampling Algorithm (EPSA) developed by [Bušić et al., 2012a] gives a rather general construction of bounding chains, under the assumption that the state space is a lattice.

ATO systems with joint returns, that motivated our work, have a very large state space (that grows exponentially with the number of item types). Furthermore the state space is *not* naturally equipped with a lattice order relation, so EPSA cannot be used. Our goal is two-fold:

- We propose a new exact simulation method, generalizing EPSA and based on aggregation. This Aggregated Envelope Perfect Sampling Algorithm (AEPSA) can be applied to reduce the state space and overcome the lack of a lattice structure.

- We provide a detailed treatment of exact simulation algorithms for ATO systems: More precisely, we use known algorithms (PSA or EPSA) for ATO systems with individual returns, and our new algorithm (AEPSA) for ATO systems with joint returns. Each type of ATO system considered is described below, and the choice of the algorithm (PSA, EPSA or AEPSA) depending on the system is summarized in Table 2.

Up to our knowledge, this is the first time perfect sampling techniques are applied to ATO systems (for an overview on ATO systems, see [Song and Zipkin, 2003]). We focus here on *continuous-review models*, with *exponential replenishment times* and *finite stock capacities*. As is common in the ATO literature, we assume at most one component of each type will be demanded for any item (the unit demand case). We consider two different options for the out-of-stock situation: A demand can be fulfilled partly (just the components that are available), referred to as a partial order service (POS); or lost fully, referred to as a total order service (TOS). Also, we distinguish between two different situations for the replenishment/return of components: components are either returned individually or jointly. We briefly mention only the possible solution techniques that are directly related to the models we consider in this paper (an overview is given in Table 1).

[Song et al., 1999] proposed an exact evaluation of ATO systems with individual returns, both for TOS and POS, by using a matrix geometric approach. This exact method, however, is computationally inefficient for larger problem instances.

[Xu, 1999], [Li and Xu, 2000], [Xu, 2002] studied the effect of correlation (for the arrival process) on a variety of system performance measures for correlated queueing systems, including ATO-POS systems. [Dayanik et al., 2003] presented several approximations and bounds on the performance of ATO-POS systems with individual returns.

The best known example of models that have joint returns of resources and TOS are telecommunication systems, or specifically loss networks ([Kelly, 1991]). In these networks, demands arrive, for example a phone call, that need several links to be simultaneously available. If all links are available, the call is connected. After the call is finished, all links are simultaneously released. When one or more of the links is not available, the call does not connect, and the demand for all links is lost. Although loss networks have a product-form solution, exactly computing the blocking

Returns	1. TOS	2. POS
A. Individual	**Exact** (matrix geometric): [Song et al., 1999]	**Exact** (matrix geometric): [Song et al., 1999] **Bounds**: [Xu, 1999], [Li and Xu, 2000], [Xu, 2002], [Dayanik et al., 2003]
B. Joint	**Exact** (product form): [Kelly, 1991]	**Approximations**: [Vliegen and van Houtum, 2009] **Bounds**: [Bušić et al., 2012b]

Table 1 Solution methods for ATO systems.

probabilities for this system is known to be a difficult problem ([Louth et al., 1994]), due to the normalizing constant.

An example of ATO systems with joint returns and a partial order service is the service tool problem, considered in [Vliegen and van Houtum, 2009]. In this problem, to perform a maintenance action, several service tools are needed at the same time. After usage, all tools return to the location they were sent from together. Whenever one or more tools are not present, they are sent by an emergency shipment to enable the initiation of the maintenance action as soon as possible. For the supply location under consideration the demand for these emergency shipped tools is lost. [Vliegen and van Houtum, 2009] developed different approximations. Some of these approximations provide provable bounds ([Bušić et al., 2012b]). For larger instances, however, these bounds are still time consuming.

Returns	1. TOS	2. POS
A. Individual	Envelope Perfect Sampling Algorithm: [Bušić et al., 2008, Bušić et al., 2012a]	Perfect Sampling Algorithm: [Propp and Wilson, 1996]
B. Joint	Aggregated Envelope Perfect Sampling Algorithm: Section 5	

Table 2 Perfect sampling algorithms for ATO systems.

For larger instances, all four cases are difficult to analyze directly and we will discuss in this paper how to develop exact simulation algorithms for ATO systems, as an alternative for bounding techniques developed in the literature. In Table 2, we mention which algorithm (PSA, EPSA or AEPSA) is used for each of the four models. The monotonicity of the POS system with individual returns (Proposition 6) allows the use of PSA, while we apply EPSA to the TOS system with individual returns, which is non-monotone (Proposition 12). To handle the case of joint returns, we use the method developed in Section 5 (AEPSA). Loss networks (case B1) have a product-form solution, which makes the exact calculation easier to some extent. We therefore focus on the other cases and we will only briefly mention in Section 7 how we can adapt the approach developed for the POS case with joint returns (Section 6) to the TOS case.

In addition, we give bounds on the complexity of our algorithms. The complexity of the exact simulation algorithms we develop depends on the coupling time that is usually difficult to estimate and even to bound, except for some specific Markov chains ([Levin et al., 2009]). In the context of queueing networks, [Dopper et al., 2006] have shown that the coupling time in an M/M/1/C queue is linear in capacity C, when $\lambda \neq \mu$ (with λ being the arrival and μ the service rate), and quadratic when $\lambda = \mu$. They used this fact to derive an upper bound for the coupling time of an acyclic network of K M/M/1/C queues that is $O(KC^2)$. These results have been extended to cyclic networks under some additional hyperstability conditions ([Anselmi and Gaujal, 2011]). However, these results use both the fact that the system is monotone under the usual product

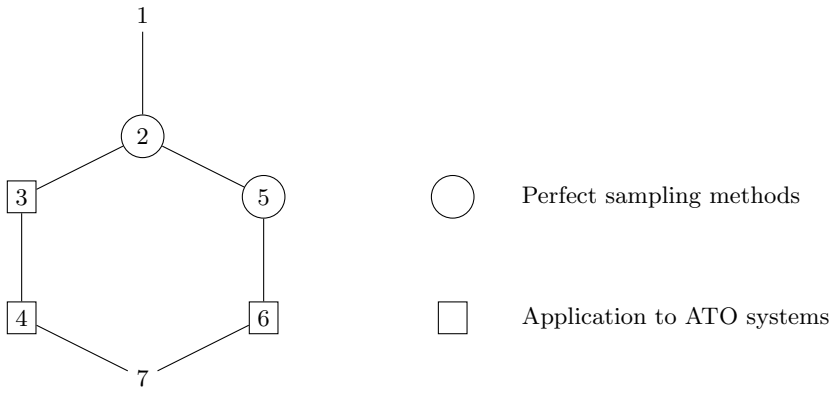


Figure 1 Dependencies between sections.

partial order and that each event in the system can only influence up to two different components. In our case, joint arrivals and services can modify many components at the same time. Also, we do not always have the monotonicity property (see TOS case).

The paper is organized as follows. In Section 2 we give an overview of perfect sampling and the related literature. Sections 3 and 4 are devoted respectively to ATO-POS and ATO-TOS systems with individual returns. In Section 5 we present AEPSA, that we apply to the ATO-POS model with joint returns in Section 6. Finally, in Section 7 we discuss some possible extensions of our work and provide conclusions. In Figure 1, we give the dependencies between sections.

2 Perfect sampling and the method of envelopes

The evolution of a finite Discrete Time Markov Chain (DTMC) can always be obtained using a finite number of discrete events (or actions). We consider a system description similar to Generalized Semi Markov Processes (see [**?**]), with a focus on state changes rather than on time: we consider a tuple $\mathcal{M} = (SS, \mathcal{E}, \nu, f)$ where SS is a finite state space, $\mathcal{E}$ is the *set of events*, ν is a probability distribution on $\mathcal{E}$, and f is a *transition function*, $f : SS \times \mathcal{E} \to SS$.

This transition function f can be naturally extended to words $a_{1\to t} \stackrel{\text{def}}{=} a_1 a_2 \dots a_t \in \mathcal{E}^t$, $t \in \mathbb{N}$ (where $a_{1\to 0} := \epsilon$ is the empty word). For any $t \in \mathbb{N}$, $f : SS \times \mathcal{E}^t \to SS$ is defined by: $f(x, \epsilon) \stackrel{\text{def}}{=} x$ and $f(x, a_{1\to t}) \stackrel{\text{def}}{=} f(f(x, a_{1\to t-1}), a_t)$ for $t \geq 1$.

Let $(a_t)_{t\geq 1}$ be an infinite *i.i.d.* sequence of random events in $\mathcal{E}$, distributed according to ν. Then for any $x_0 \in SS$, the random process $(X_t \stackrel{\text{def}}{=} f(x_0, a_{1\to t}))_{t\geq 0}$ is a Markov chain started in x_0 with probability transition matrix P given by:

$$\text{for all } x, y \text{ in } SS, \quad P(x,y) = \sum_{a\in\mathcal{E},\ f(x,a)=y} \nu(a). \tag{1}$$

We say that the Markov chain (X_t) is *generated* by $\mathcal{M}$ and $(a_t)_{t\geq 1}$.

Conversely, for any probability transition matrix P on a finite state space SS, it is easy to see that there exists a tuple $\mathcal{M} = (SS, \mathcal{E}, \nu, f)$ such that (1) holds, *i.e.* such that $\mathcal{M}$ generates a Markov chain on SS with transition matrix P, but that representation is in general not unique. However, such a representation naturally arises for many systems, including Markovian queueing networks.

We can build a family of Markov chains $\{(X_t(x) = f(x, a_{1\to t}))_{t\geq 0} \mid x \in SS\}$ starting from each state $x \in SS$, referred to as the *grand coupling* generated by $\mathcal{M}$ and $(a_t)_{t\geq 1}$ ([Levin et al., 2009]). We will say that the grand coupling has *coupled* (or more precisely *coalesced*) at time t if all the Markov chains of the family have reached the same state. Using the notation $f(U, a_{1\to t}) \stackrel{\text{def}}{=}$

$\{f(x, a_{1\to t}), x \in U\}$ for any subset $U \subset SS$, this is equivalent to the fact that $f(SS, a_{1\to t})$ is reduced to a singleton. In the following, $|V|$ denotes the cardinality of set V.

2.1 Perfect Sampling

Let $(X_t)_{t\in\mathbb{N}}$ be an irreducible and aperiodic DTMC with finite state space SS and transition matrix P. Consider a discrete event system representation $\mathcal{M} = (SS, \mathcal{E}, \nu, f)$ that satisfies (1), and let π denote the steady state distribution of the chain. Perfect Sampling Algorithm (PSA) gives a sample from the steady state distribution in finite time, using a coupling from the past construction.

Theorem 1 ([Propp and Wilson, 1996])**.** *Let $(a_{-t})_{t\in\mathbb{N}} = (a_0, a_{-1}, \dots, a_{-t}, \dots)$ be a sequence of* i.i.d. *events with distribution ν on $\mathcal{E}$. There exists $\ell \in \mathbb{N}$ such that $\lim_{t\to\infty} \left|f(SS, a_{-t+1\to 0})\right| = \ell$ almost surely. The grand coupling generated by $\mathcal{M}$ and $(a_{-t})_{t\in\mathbb{N}}$ is coalescing if $\ell = 1$. In that case, let*

$$\tau \stackrel{\text{def}}{=} \inf\left\{t \,:\, \left|f(SS, a_{-t+1\to 0})\right| = 1\right\}$$

be the coupling time of the chain. Then $\mathbb{E}(\tau) < \infty$ and $f(SS, a_{-\tau+1\to 0})$ is steady state distributed.

The main drawback of PSA is the fact that one needs to simulate one Markov chain starting from each state in $\mathcal{X}$, which is too large for most applications. Several approaches have been used to overcome this problem. The main one for a partially ordered state space $(\mathcal{X}, \preceq)$ and monotone events was already given in [Propp and Wilson, 1996].

Definition 2. *An event $a \in \mathcal{E}$ is said to be monotone if, for all $x, y \in \mathcal{X}$, $x \preceq y \;\Rightarrow f(x, a) \preceq f(y, a)$.*

If all events are monotone, then one can consider only the trajectories issued from maximal and minimal initial states ([Propp and Wilson, 1996]). In the case of general non-monotone chains, it is possible to use a bounding chain method, introduced in [Kendall and Møller, 2000]. EPSA (Envelope Perfect Sampling Algorithm, [Bušić et al., 2008, Bušić et al., 2012a]) constructs bounding chains in the case when the state space is equipped with a lattice order relation. We give next a short overview of EPSA.

2.2 Bounding Interval Chains

Let $(\mathcal{X}, \preceq)$ be a lattice. For $m, M \in \mathcal{X}$, denote by $[m, M] \stackrel{\text{def}}{=} \{x \in \mathcal{X} \,:\, m \preceq x \preceq M\}$ the *(lattice) interval* between the endpoints m and M. Let $\mathcal{J}$ be the set of all nonempty lattice intervals: $\mathcal{J} = \{[m, M] \,:\, m, M \in \mathcal{X},\, m \preceq M\}$. Given a grand coupling $\{(X_t(x))_{t\geq 0} \mid x \in \mathcal{X}\}$, a *bounding interval chain* is any Markov chain of nonempty intervals $([m_t, M_t])_{t\geq 0}$ such that: for all x in $\mathcal{X}$ and all $t \geq 0$, $X_t(x) \in [m_t, M_t]$. In particular we notice that when $m_t = M_t$, the grand coupling has necessarily coalesced.

An *envelope transition* function $F : \mathcal{J} \times \mathcal{E} \to \mathcal{J}$ is defined by: for all $[m, M] \in \mathcal{J}$ and $a \in \mathcal{E}$,

$$F([m, M], a) \stackrel{\text{def}}{=} \left[\inf_{m\preceq x\preceq M} f(x, a), \sup_{m\preceq x\preceq M} f(x, a)\right]. \tag{2}$$

As with f, the transition function F can be extended to finite words of events. For any $t \in \mathbb{N}$, $F : \mathcal{J} \times \mathcal{E}^t \to \mathcal{J}$ is defined by: $F([m, M], \epsilon) \stackrel{\text{def}}{=} [m, M]$ and $F([m, M], a_{1\to t}) \stackrel{\text{def}}{=} F(F([m, M], a_{1\to t-1}), a_t)$ for $t \geq 1$.

Let $\bot \stackrel{\text{def}}{=} \inf \mathcal{X}$ (resp . $\top \stackrel{\text{def}}{=} \sup \mathcal{X}$) be the bottom (resp. top) element of $\mathcal{X}$. The process $[m_t, M_t] \stackrel{\text{def}}{=} F([\bot, \top], a_{1\to t})$ is a Markov chain over the state space $\mathcal{X} \times \mathcal{X}$, called the *envelope chain*, and is a bounding interval chain of the grand coupling $\{(f(x, a_{1\to t}))_{t\geq 0} \mid x \in \mathcal{X}\}$.

The envelope process can be used to detect the coalescence of the grand coupling. The following result was shown in [Bušić et al., 2008]:

Theorem 3. *Let $(a_{-t})_{t\in\mathbb{N}}$ be a sequence of* i.i.d. *events with distribution ν on $\mathcal{E}$. Assume that the envelope chain $F([\bot,\top],a_{-t+1\to 0})$ hits the set of single point intervals $\mathcal{P}=\{[x,x]\,:\,x\in\mathcal{X}\}$ a.s. in finite time. Let $\tau_e \stackrel{\text{def}}{=} \min\{t\,:\,F([\bot,\top],a_{-t+1\to 0})\in\mathcal{P}\}$, then τ_e is a backward coupling time of the envelope chain. The state defined by $F([\bot,\top],a_{-\tau_e+1\to 0})$ has the steady state distribution of DTMC $(X_t)_{t\in\mathbb{N}}$.*

Algorithm 1: Envelope Perfect Sampling Algorithm (EPSA)

Data: *I.i.d.* events $(a_{-t})_{t\in\mathbb{N}}\in\mathcal{E}^{\mathbb{N}}$
Result: A state $x^*\in\mathcal{X}$ generated according to the stationary distribution of the Markov chain

```
begin
    t := 1;
    repeat
        m := ⊥; M := ⊤;
        for i = t − 1 downto 0 do
            [m, M] := F([m, M], a_{−i}) ;
        t := 2t;
    until m = M;
    x* := m;
    return x*;
```

Envelope Perfect Sampling Algorithm (EPSA) is given in Algorithm 1. The reason to double t at each iteration of the algorithm is that we need to compute $F([\bot,\top],a_{-t+1\to 0})$ in each loop, which corresponds to t iterations of F. While increasing t by 1 would lead to a quadratic cost in τ_e, doubling it keeps the complexity linear. This was already observed in [Propp and Wilson, 1996], for the monotone case.

The construction of the envelope chain depends on the discrete event representation of the Markov chain that is not unique. Different event representations lead to different envelope chains with different coupling properties (one may coalesce almost surely and the other not, or if they both coalesce their coupling times may be different). The complexity of the envelope transition function may also differ depending on the representation.

Remark 4. *When the assumptions of Theorem 3 do not hold (i.e. the envelope chain does not couple), EPSA never stops. However in that case, the variants of EPSA can still provide perfect samples or performance bounds for increasing cost functions:*

- *EPSA with splitting, proposed by [Bušić et al., 2008], can still generate perfect samples in finite time, under milder assumptions. The splitting algorithm is hybrid: it first runs EPSA when the envelopes are too far apart and switches to the usual PSA algorithm as soon as the number of states inside the envelopes becomes manageable.*

- *If we stop the algorithm at any time (for instance after some maximal number of iterations is reached), we will obtain an interval estimate for the stationary distribution. Indeed, let $X\sim\pi$ be a random variable distributed according to the stationary distribution π. For any $s\geq 0$, $[\bot_s,\top_s]\stackrel{\text{def}}{=}F([\bot,\top],a_{-s+1\to 0})$ satisfies:*

$$\bot_s \preceq_{st} X \preceq_{st} \top_s, \tag{3}$$

 where $\preceq_{st}$ denotes the usual strong stochastic order of random variables. For two random variables X and Y with values in $(\mathcal{X},\preceq)$, $X\preceq_{st}Y$ if $P(X\in U)\leq P(Y\in U)$ for all increasing sets $U\subset\mathcal{X}$ (see [?] for further material on stochastic orderings).

 This interval estimate can be used to obtain performance bounds, as explained in the following subsection.

2.3 Performance Evaluation Bounds

Assume now $c : \mathcal{X} \to \mathbb{R}$ is some increasing cost function defined on states of the Markov chain (e.g. the total number of items in replenishment, or the probability that a new demand cannot be fulfilled). Then (3) implies:

$$\mathbb{E}[c(\bot_s)] \leq \mathbb{E}[c(X)] \leq \mathbb{E}[c(\top_s)],\ s \geq 0,$$

so we can use EPSA to obtain bounds of the steady-state cost. In performance evaluation of a given ATO system, or comparison between two different system designs, we are often interested in some specific performance guarantees (e.g. probability that a new demand cannot be fulfilled must be below a certain level). In that case, we are only interested in bounds for a given cost function and not the exact samples from the stationary distribution. Furthermore, the difference between the lower and upper bound provides also the error estimate for these bounds.

In Sections 3 and 4 we apply these existing methods (PSA and EPSA) to ATO systems with individual replenishments (PSA for the POS case in Section 3 and EPSA for the TOS case in Section 4).

3 ATO-POS with individual state-dependent replenishments

3.1 Model description

We consider the ATO-POS system with individual replenishments of items. There are I different item types and let $\mathcal{I} = \{1, \ldots, I\}$. We assume finite stock capacities and denote by C_i the total amount of items of type $i \in \mathcal{I}$. Customers arrive in the system according to a Poisson process of rate λ. Each customer asks for a subset of items and the probability to ask for subset A is denoted by p_A, *i.e.* the demands for each subset A follow a Poisson process of rate $\lambda_A = p_A\lambda$. If some demanded items are not available, then the customer takes the available items (POS case) and the demand for the items that are not available is lost. As often considered in the ATO literature, we assume that the number of different subsets customers can ask for is small (for instance, $|\{A \subset \mathcal{I} : \lambda_A \neq 0\}|$ is linear with respect to the number I of item types). Each item of type i is replenished after an exponential time, with a rate that depends on the current amount of items i in replenishment. We assume that the replenishments of different item types are mutually independent, and independent from the demands.

This system can be modeled as a network of I queues with joint arrivals and independent services: arrivals to queues represent demands for different subsets of items and services in a queue model replenishments of items. Denote by $C = (C_1, \ldots, C_I)$ the vector of queue capacities.

The total number of items in each queue (i.e. in replenishment) is given by a vector $x = (x_1, \ldots, x_I)$, where x_i is the number of items of type i. The state space of the system is: $\mathcal{X} = \{0, ..., C_1\} \times \{0, ..., C_2\} \times \cdots \times \{0, ..., C_I\}$. In the following, for $i \in \mathcal{I}$, we denote by e_i the state with all the components equal to 0, except component i that is equal to 1.

We have two different types of transitions. For each $x \in \mathcal{X}$, $A \subset \mathcal{I}$ and for each $i \in \mathcal{I}$:

- There is a demand for subset A, with rate λ_A. The new state is: $x + \sum_{j \in A} \mathbb{1}_{\{x_j < C_j\}} e_j$.
- If $x_i > 0$, there is a service in queue i, with rate $\mu_i(x_i)$ that depends on the current number x_i of items of type i. The new state is $x - e_i$.

By a standard uniformization procedure, we can transform this continuous time Markov chain to a discrete time Markov chain. Let $\beta_i := \max_{1 \leq x_i \leq C_i} \mu_i(x_i)$ be the maximal service rate for queue i. Then the outgoing rate in each state is upper-bounded by $\Lambda := \lambda + \sum_{i \in \mathcal{I}} \beta_i$. We take the uniformization constant equal to Λ.

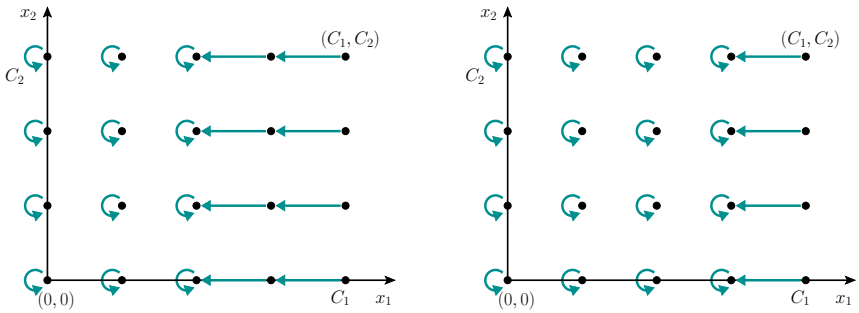


Figure 2 Case $\mu_1(x_1) = \mu_1 \cdot x_1$. On the left: Transition function for service $s_1^{(3)}$. On the right: Transition function for service $s_1^{(4)}$.

Event representation.

We now explain a discrete event representation of our (uniformized) Markov chain. In order to allow the construction of a family of Markov chains on the same probability space and driven by the same sequence of events, the set of possible events and the event rates cannot depend explicitely on the state. The following events allow such a construction.

- *Arrivals.* For any $A \subset \mathcal{I}$, $A \neq \emptyset$, let d_A be the event of probability λ_A/Λ that corresponds to a "joint arrival to queues in A".
- *Services.* As the service rate in each queue does not depend on the state of other queues, we can consider the queues separately. For queue i, the service rate is given by a function $\mu_i(x_i)$. We start by reordering the set $\{0, \ldots, C_i\}$ of possible values of x_i in increasing order of μ_i, and denote this permutation by $\ell_i = (\ell_i^{(0)}, ..., \ell_i^{(C_i)})$. We have:

$$0 = \mu_i(\ell_i^{(0)}) \leq \mu_i(\ell_i^{(1)}) \leq ... \leq \mu_i(\ell_i^{(C_i)}) = \beta_i.$$

 For each $1 \leq j \leq C_i$, let $s_i^{(j)}$ be the event of probability $\left[\mu_i(\ell_i^{(j)}) - \mu_i(\ell_i^{(j-1)})\right]/\Lambda$ that corresponds to a "service in queue i for states $x = (x_1, ..., x_I)$ such that $x_i \in \{\ell_i^{(j)}, \ell_i^{(j+1)}, ..., \ell_i^{(C_i)}\}$".

Example 5. *We assume that $\mu_i(x_i) = \mu_i \cdot x_i$ for some $i \in \mathcal{I}$, and consider the events for services in queue i. For $1 \leq j \leq C_i$, we define the event $s_i^{(j)}$ of probability μ_i/Λ, as a service in the i-th queue for all the states x such that $x_i \geq j$. In particular, states x such that $x_i = C_i$ (that have the highest service rate in queue i) are served in each of these events, while, for states x such that $x_i = 1$, the number of items in the i-th queue decreases only when event $s_i^{(1)}$ occurs. For $I = 2$, $i = 1$ and $C_1 = 4$, the transition function for events $s_1^{(3)}$ and $s_1^{(4)}$ is given in Figure 2.*

The set of events in the system is: $\mathcal{E} = \{d_A, A \subset \mathcal{I}, A \neq \emptyset\} \cup \{s_i^{(j)}, i \in \mathcal{I}, 1 \leq j \leq C_i\}$ and the probability distribution ν on $\mathcal{E}$ is given by: $\nu(d_A) = \lambda_A/\Lambda$, $\emptyset \neq A \subset \mathcal{I}$, $A \neq \emptyset$, and $\nu(s_i^{(j)}) = \left[\mu_i(\ell_i^{(j)}) - \mu_i(\ell_i^{(j-1)})\right]/\Lambda$, $i \in \mathcal{I}, 1 \leq j \leq C_i$. Note that some events may have probability 0; we can ignore these events (*i.e.* restrict $\mathcal{E}$ to the support of ν).

Transition function.

The transition function $f_P : \mathcal{X} \times \mathcal{E} \to \mathcal{X}$ of the ATO-POS system is defined as follows. For $x \in \mathcal{X}$, $A \subset \mathcal{I}$, $A \neq \emptyset$, $i \in \mathcal{I}$ and $1 \leq j \leq C_i$:

$$\begin{cases} f_P(x, d_A) &= x + \sum_{k \in A} \mathbb{1}_{\{x_k < C_k\}} e_k, \\ f_P(x, s_i^{(j)}) &= x - \sum_{k=j}^{C_i} \mathbb{1}_{\{x_i = \ell_i^{(k)}\}} e_i \;=\; x - \mathbb{1}_{\left\{\mu_i(x_i) \geq \mu_i\left(\ell_i^{(j)}\right)\right\}} e_i. \end{cases}$$

Monotonicity.

We consider the natural product order on $\mathcal{X}$, *i.e.* $x = (x_1, ..., x_I) \leq (y_1, ..., y_I) = y$ if $x_i \leq y_i$ for all $1 \leq i \leq I$. By using the fact that a service $s_i^{(j)}$ can only modify the i-th coordinate, it easily follows that:

Proposition 6. *The transition function f_P of the ATO-POS system is monotone under the natural product order on $\mathcal{X}$, i.e. for all $a \in \mathcal{E}$, $x \leq y \implies f_P(x, a) \leq f_P(y, a)$.*

Since the system is monotone, it is enough to consider only two trajectories (starting from the upper and lower states). The complexity of the Perfect Sampling Algorithm depends on the coupling time of these two trajectories.

3.2 Bound for the coupling time

Let $(a_{-t})_{t \in \mathbb{N}} = (a_0, a_{-1}, \ldots, a_{-t}, \ldots)$ be a sequence of *i.i.d.* events with distribution ν on $\mathcal{E}$. Let $\tau^P := \min\left\{t \;:\; \left|f_P(\mathcal{X}, a_{-t+1 \to 0})\right| = 1\right\}$ be the coupling time of the whole chain X (defined on $\mathcal{X}$). For all $1 \leq i \leq I$, we define the projection on i-th component as:

$$\phi_i : \left\{ \begin{array}{ccc} \mathcal{X} & \longrightarrow & [0, C_i] \\ x = (x_1, \ldots, x_I) & \longmapsto & x_i \end{array} \right. ,$$

and set $\tau^{P,i} := \min\left\{t \;:\; \left|\phi_i\left(f_P(\mathcal{X}, a_{-t+1 \to 0})\right)\right| = 1\right\}$, the "coupling time" of the chain on the i-th component.

We first prove the following property:

$$\left(\left|\phi_i\left(f_P(\mathcal{X}, a_{-s+1 \to 0})\right)\right| = 1 \text{ and } t \geq s \right) \quad \Longrightarrow \quad \left|\phi_i\left(f_P(\mathcal{X}, a_{-t+1 \to 0})\right)\right| = 1. \tag{4}$$

Informally, we say that, as soon as a component couples, it *stays coupled*. This property implies that $\tau^P = \max_{1 \leq i \leq I} \tau^{P,i}$.

Definition 7. *A Markov chain that satisfies property (4) for any sequence of events $(a_{-t})_{t \in \mathbb{N}}$ is said to be componentwise coupling. In that case, the whole coupling time of the chain can be expressed in terms of the coupling times on each component.*

The following lemma shows that the chain X is componentwise coupling.

Lemma 8. *Let $i \in \mathcal{I}$. The projection $\phi_i(X)$ of the chain on the i-th component is a Markov chain on the state space $[0, C_i]$. Let $f_{P,i}$ be its transition function. Then we have that:*

$$\tau^{P,i} = \min\left\{t \;:\; \left|f_{P,i}([0, C_i], a_{-t+1 \to 0})\right| = 1\right\},$$

i.e. $\tau^{P,i}$ is the coupling time of the Markov chain $\phi_i(X)$ (it does not depend on the value of the chain X on other components).

Proof. We compute $\phi_i \circ f_P$. For $x \in \mathcal{X}$, $A \subset \mathcal{I}$, $A \neq \emptyset$, $i \in \mathcal{I}$ and $1 \leq j \leq C_i$:

$$\left\{ \begin{array}{rcl} \phi_i \left(f_P(x, d_A) \right) & = & x_i + \mathbb{1}_{\{i \in A \ \& \ x_i < C_i\}}, \\ \phi_i \left(f_P(x, s_i^{(j)}) \right) & = & x_i - \mathbb{1}_{\left\{ \mu_i(x_i) \geq \mu_i\left(\ell_i^{(j)}\right) \right\}}. \end{array} \right.$$

Hence we have that $\phi_i \circ f_P$ only depends on x_i. For any event $a \in \mathcal{E}$, we can set $f_{P,i}(x_i, a) = \phi_i \left(f_P(x, a) \right)$. This ends the proof. □

As a direct application of Lemma 8, we have:

$$\mathbb{E}[\tau^P] = \mathbb{E}\left[\max_{1 \leq i \leq I} \tau^{P,i} \right] \leq \sum_{i=1}^{I} \mathbb{E}[\tau^{P,i}]. \tag{5}$$

In order to bound $\mathbb{E}[\tau^P]$, we will show bounds on $\mathbb{E}[\tau^{P,i}]$. We use the following result that is often used as a part of the proof of Foster's stability criterion (see for instance [**?**, proof of Theorem 1.1]).

Lemma 9. *Let the transition matrix P on the finite state space S be irreducible and suppose that there exists a function $h : S \to \mathbb{R}_+$ such that*

$$\sum_{z \in S} P(y, z) h(z) \leq h(y) - \epsilon \text{ for all } y \notin U, \tag{6}$$

for some subset $U \subset S$. Let τ_U be the hitting time of U and $\mathbb{E}_y$ denote the expectation, knowing that the chain starts in y. Then, for all $y \notin U$,

$$\mathbb{E}_y[\tau_U] \leq \frac{h(y)}{\epsilon}. \tag{7}$$

We next give a bound for $\mathbb{E}(\tau^{P,i})$ for the two following cases: either service or arrival rate is high.

Lemma 10. *Let $i \in \mathcal{I}$. Let $\lambda_i := \sum_{A \ni i} \lambda_A$ be the total arrival rate in queue i. Set $\alpha_i := \min_{1 \leq x_i \leq C_i} \mu_i(x_i)$ and $\eta_i := \max_{1 \leq x_i < C_i} \mu_i(x_i)$. Then:*

- *If $\lambda_i < \alpha_i$, then $\mathbb{E}[\tau^{P,i}] \leq \frac{\Lambda C_i}{\alpha_i - \lambda_i}$.*
- *If $\lambda_i > \eta_i$, then $\mathbb{E}[\tau^{P,i}] \leq \frac{\Lambda C_i}{\lambda_i - \eta_i}$.*

Proof. Let $S = \{0, ..., C_i\}$ and P be the transition matrix of the Markov chain $Y = \phi_i(X)$. When the chain Y is in state y, $0 \leq y \leq C_i$: Y goes to $y + \mathbb{1}\{y < C_i\}$ with probability λ_i / Λ, it goes to $y - 1$ with probability $\mu_i(y)/\Lambda$, and stays at y with probability $\left[\Lambda - \lambda_i - \mu_i(y) \right] / \Lambda$.

In the first case ($\lambda_i < \alpha_i$), we use Lemma 9 with $U = \{0\}$, $h(z) = z$ for all $z \in S$, $\epsilon = (\alpha_i - \lambda_i)/\Lambda$ and apply (7) with $y = C_i$. This gives that $\mathbb{E}[\tau^{P,i}] \leq \mathbb{E}_{C_i}[\tau_{\{0\}}] \leq \frac{h(C_i)}{\epsilon} = \frac{\Lambda C_i}{\alpha_i - \lambda_i}$, where the first inequality comes from the fact that the chain X (and thus its projection $\phi_i(X)$) is monotone (Proposition 6): when the chain starting from $y = C_i$ reaches 0, the chains starting from all other states are in 0, and the system coupled. In the second case ($\lambda_i > \eta_i$), we use Lemma 9 with $U = \{C_i\}$, $h(z) = C_i - z$ for all $z \in S$, $\epsilon = (\lambda_i - \eta_i)/\Lambda$ and apply (7) with $y = 0$. □

This lemma and equation (5) lead to the following proposition:

Proposition 11. *Let $\mathcal{I}_0 = \{i : 1 \leq i \leq I, \lambda_i < \alpha_i\}$, $\mathcal{I}_C = \{i : 1 \leq i \leq I, \lambda_i > \eta_i\}$ and assume $\mathcal{I} = \mathcal{I}_0 \cup \mathcal{I}_C$. Then we have:*

$$\mathbb{E}[\tau^P] \leq \Lambda \left(\sum_{i \in \mathcal{I}_0} \frac{C_i}{\alpha_i - \lambda_i} + \sum_{i \in \mathcal{I}_C} \frac{C_i}{\lambda_i - \eta_i} \right).$$

We discuss two important cases:

- *Single server case.* Assume $\mu_i(x_i) = \mu_i$ for all $i \in \mathcal{I}$, $x \in \mathcal{X}$. Then $\alpha_i = \eta_i = \mu_i$, and $\Lambda = \lambda + \sum_i \mu_i$. Hence the mean coupling time of the chain X is in $O(|C|)$, with $|C| = \sum_i C_i$. In addition we can notice that the hypothesis $\mathcal{I} = \mathcal{I}_0 \cup \mathcal{I}_C$ is necessary to have a linear bound in $|C|$: $\mathcal{I} \neq \mathcal{I}_0 \cup \mathcal{I}_C$ implies $\lambda_i = \mu_i$ for some $i \in \mathcal{I}$, and $\mathbb{E}[\tau^{P,i}]$ is quadratic in C_i in that case (see [Dopper et al., 2006]).
- *Infinite server case.* Assume $\mu_i(x_i) = \mu_i \cdot x_i$ for all $i \in \mathcal{I}$, $x \in \mathcal{X}$. Then $\alpha_i = \mu_i$, $\eta_i = \mu_i \cdot (C_i - 1)$ and $\Lambda = \lambda + \sum_i \mu_i \cdot C_i$. If $\mathcal{I} = \mathcal{I}_0$, the mean coupling time of the chain is in $O\left(|C|^2\right)$. Note that this bound is larger than in the single server case as a result of the time-discretization (uniformization) of our chain.

In Appendix A, we consider a slightly more general model: the service rates μ_i can depend on other components, but the service events remain monotone. Under a high service rate assumption, we give a bound on the mean hitting time to zero (that provides an upper bound for the coupling time). This result will be also used to bound the running time of the AEPSA algorithm for ATO-POS with joint returns, studied in Section 6.

4 ATO-TOS with individual state-dependent replenishments

4.1 Model description

We consider the ATO-TOS model with individual replenishments of items. The difference with the ATO-POS model is in the way the demands are handled in the out-of-stock situation: If some demanded items are not available, then the whole demand is lost.

As before, we model the system by a queueing system, with state space $\mathcal{X}$. We consider the uniformized Markov chain with uniformization constant $\Lambda = \lambda + \sum_{i \in \mathcal{I}} \beta_i$, with $\beta_i = \max_{1 \leq x_i \leq C_i} \mu_i(x_i)$ as before. We consider the same set $\mathcal{E}$ of events and the same probability distribution ν on $\mathcal{E}$ as in the POS case (defined in Section 3.1). The transition function for services is also the same as in the POS case. For an arrival d_A, $A \subset \mathcal{I}$, $A \neq \emptyset$, the response of POS and TOS systems is different only for states x that belong to the boundary of the state space (*i.e.* such that there exists $i \in A$ with $x_i = C_i$).

Transition function.

The transition function f_T for the ATO-TOS system is defined as follows. Let $A \subset \mathcal{I}$, $A \neq \emptyset$, $i \in \mathcal{I}$, $1 \leq j \leq C_i$, and $x \in \mathcal{X}$:

$$\begin{cases} f_T(x, d_A) &= x + \left(\prod_{i \in A} \mathbb{1}_{\{x_i < C_i\}} \right) e_A, \\ f_T(x, s_i^{(j)}) &= f_P(x, s_i^{(j)}), \end{cases}$$

where $e_A = \sum_{i \in A} e_i$. In other words, when an arrival d_A occurs, we add 1 to each component of x in A, if *all* its components x_i, $i \in A$, satisfy $x_i < C_i$.

In the ATO-TOS system, whether or not an arrival in queue i is accepted depends on the whole state of the system, which makes the system more difficult to study. Indeed, contrary to ATO-POS, arrivals are not necessarily monotone in the ATO-TOS system.

4.2 Envelopes

(Non-)monotonicity.

As before, we can consider the product order on $\mathcal{X}$. Since services are the same as in the ATO-POS model, they are monotone (Proposition 6). For the same reason, the arrivals of *only one* item are also monotone. Unfortunately, as soon as $|A| \geq 2$, the event d_A is not monotone for the product

order on $\mathcal{X}$. For instance, let $I = 2$, $x = (C_1 - 1, C_2 - 1)$ and $y = (C_1 - 1, C_2)$. Then $x \leq y$, yet $f_T(x, d_{\{1,2\}}) = (C_1, C_2) \geq y = f_T(y, d_{\{1,2\}})$.

One could try to find another partial order on $\mathcal{X}$ for which the ATO-TOS model would be monotone. If $I = 2$, it is easy to check that this is true for the following partial order: $x = (x_1, x_2) \preceq (y_1, y_2) = y$ if $x_1 \geq y_1$ and $x_2 \leq y_2$. However, for $I \geq 3$ we show much stronger statement (the proof is given in Appendix B):

Proposition 12. *Let us consider the ATO-TOS system with $I \geq 3$. If f_T is monotone for a partial order $\preceq$, then $\preceq$ is the trivial order, i.e. : $x \preceq y \Leftrightarrow x = y$.*

We now show how to appy EPSA (Algorithm 1) to the ATO-TOS model.

Computation of envelopes.

We consider the product order on $\mathcal{X}$, and define the envelope transition function F_T as in (2), *i.e.* for all m, $M \in \mathcal{X}$ such that $m \leq M$, and $a \in \mathcal{E}$, we set:

$$F_T([m, M], a) \stackrel{\text{def}}{=} \left[\inf_{m \leq x \leq M} f_T(x, a), \sup_{m \leq x \leq M} f_T(x, a) \right].$$

Services and arrivals of only one object are monotone for f_T, therefore envelopes follow easily: let m, $M \in \mathcal{X}$ such that $m \leq M$, and consider a monotone event a, then we have $F_T([m, M], a) = [f_T(m, a), f_T(M, a)]$.

Thus we are left with the computation of envelopes for an arrival d_A of several objects. As a consequence of the following proposition, we have that the computation of envelopes can be done in a linear time with respect to the number I of queues.

Proposition 13. *Let $A \subset \mathcal{I}$ such that $|A| \geq 2$. Let m, $M \in \mathcal{X}$, $m \leq M$. We distinguish three cases:*

Case I: *For all $i \in A$, $M_i < C_i$. Then $F_T([m, M], d_A) = [m + e_A, M + e_A]$.*

Case II: *There exists $i_0 \in A$ such that $m_{i_0} = C_{i_0}$. Then $F_T([m, M], d_A) = [m, M]$.*

Case III: *Otherwise, i.e. if there exists $i_0 \in A$ such that $M_{i_0} = C_{i_0}$ and for all $i \in A$, $m_i < C_i$. Let $[m', M'] := F_T([m, M], d_A)$. Then:*

$$\begin{aligned} m'_{i_0} &= m_{i_0} + \prod_{k \in A \setminus \{i_0\}} \mathbb{1}_{\{M_k < C_k\}}, \\ m'_i &= m_i, \; i \neq i_0, \end{aligned}$$

and $M'_i = M_i + \mathbb{1}_{\{i \in A,\, M_i < C_i\}}$.

Proof. Cases I and II are straightforward. In case I, for all $x \in \mathcal{X}$ such that $m \leq x \leq M$, we have $f_T(x, d_A) = x + e_A$. In case II, for all $x \in \mathcal{X}$ such that $m \leq x \leq M$, we have $f_T(x, d_A) = x$.

Case III. If $i \notin A$, then $m'_i = m_i$ and $M'_i = M_i$. Let $i \in A$.

The upper envelope is simpler. Clearly, $M'_i \leq M_i + \mathbb{1}_{\{M_i < C_i\}}$, so we only need to find x, $m \leq x \leq M$, such that $(f_T(x, d_A))_i = M_i + \mathbb{1}_{\{M_i < C_i\}}$. This is true for state x such that $x_i = M_i$ and $x_j = m_j$ for all $j \neq i$ (by assumption, for all $j \in A$, $m_j < C_j$, so an arrival d_A cannot be 'blocked' by $j \neq i$).

Now we compute the lower envelope. Clearly, $m'_i \geq m_i$, for all i. For state x such that $x_{i_0} = M_{i_0}$ and $x_i = m_i$, $i \neq i_0$, we get $(f_T(x, d_A))_i = m_i$, $i \neq i_0$, so $m'_i = m_i$, $i \neq i_0$. For component i_0, we distinguish two cases:

1. $\exists k \in A \setminus \{i_0\}$, $M_k = C_k$. This case is similar as before, as we can take $x_k = M_k$ and $x_i = m_i$, $i \neq k$ (x is 'blocked' by the k-th component, and $f_T(x, d_A) = x$). As $i_0 \neq k$, we get $m'_{i_0} = m_{i_0}$.

2. $\forall k \in A \setminus \{i_0\}$, $M_k < C_k$. For all $x \in \mathcal{X}$ such that $x_{i_0} = m_{i_0}$, we have that $x_k \leq M_k < C_k$ for $k \in A \setminus \{i_0\}$, so $f_T(x, d_A) = x + e_A$. Hence $(f_T(x, d_A))_{i_0} = m_{i_0} + 1$.

□

4.3 Bound for the coupling time

Let τ^T be the coupling time of the ATO-TOS system. Using a coupling between POS and TOS models and a bound on the mean hitting time to zero for the POS model give the following proposition:

Proposition 14. *Assume* $\delta = \min_{x \neq 0} \sum_i \mu_i(x) - \sum_i \lambda_i > 0$. *Then we have:* $\mathbb{E}[\tau^T] \leq \frac{\Lambda}{\delta}|C|$, *where* $|C| = \sum_{i=1}^{I} C_i$.

Before proving this proposition, we remark that, as for Proposition 11, the mean coupling time of the chain X is in $O(|C|)$ for the single server case ($\mu_i(x_i) = \mu_i$ for all $i \in \mathcal{I}$), and in $O\left(|C|^2\right)$ when $\mu_i(x_i) = \mu_i \cdot x_i$ for all $i \in \mathcal{I}$ and $x \in \mathcal{X}$.

In order to prove Proposition 14, we first need the following lemma:

Lemma 15. *Let* $a \in \mathcal{E}$ *and* $x, y \in \mathcal{X}$. *Then* $x \leq y \;\Rightarrow\; f_T(x, a) \leq f_P(y, a)$.

Proof of Lemma 15. Result for services follows from the fact that they are the same in both of the models, and from their monotonicity. In the case of an arrival d_A, $A \subset \mathcal{I}$, $A \neq \emptyset$, we have, for all $x \in \mathcal{X}$: $f_T(x, d_A) \leq f_P(x, d_A)$. We can conclude using the monotonicity of arrivals for the ATO-POS model. □

Proof of Proposition 14. Let τ_0^P be the hitting time to zero for the ATO-POS model. We show that

$$\tau^T \leq \tau_0^P. \tag{8}$$

Then Proposition 24 (Appendix A) gives a bound on the mean of τ_0^P and concludes the proof. Now we prove (8). Assume we do coupling from the past, starting from state C, for both ATO-POS and ATO-TOS (with a coupling using the same events). When ATO-POS is in state zero, ATO-TOS is also in state zero by Lemma 15, and hence it has coupled. So (8) follows.

Note that we cannot bound τ^T by the coupling time τ^P of ATO-POS (see Figure 3 and the following subsection). □

4.4 Comparison between POS and TOS models

In Figure 3, on the left, we give coupling times for ATO-POS and ATO-TOS models with individual services (Sections 3 and 4), for the following parameters: $I = 5$, $C_i = 10, \forall i$, $\lambda_A = \frac{1}{2^{|A|-1}}$, and $\mu_i(x_i) = \mu x_i$, with $\rho = \frac{\lambda_i}{\mu}$. The size of the sample is $n = 100$. We can see that the upper bounds for ATO-POS (see Proposition 11) are quite pessimistic for this example, in particular when $\lambda_i < \alpha_i$. This is mainly due to the fact that these bounds were obtained by considering the sum of the mean coupling times for each dimension, while the coupling of other components may occur faster knowing that one component has already coupled. In addition, when $\lambda_i >> \mu$, we can observe that the bound we obtained for the POS model is not a bound for the TOS model (even if the transition function of the POS model is a bound for the one of the TOS model). This is an illustration of the fact that the chain in the TOS model is not componentwise coupling (see the definition above Lemma 8, in 3.2). This explains why we use the hitting time to zero in the POS system, in order to bound the coupling time of the TOS chain.

In Figure 3, on the right, we give the mean values of the total number of jobs in the system (*i.e.* $\sum_i x_i$) for ATO-POS and ATO-TOS, for the same parameter values. One can see that both models are very close when $\rho < 1$, and the difference between the two models only become significant when $\lambda_i >> \mu$, which is expected as the difference between them is only for joint arrivals when the system reaches its capacity limits.

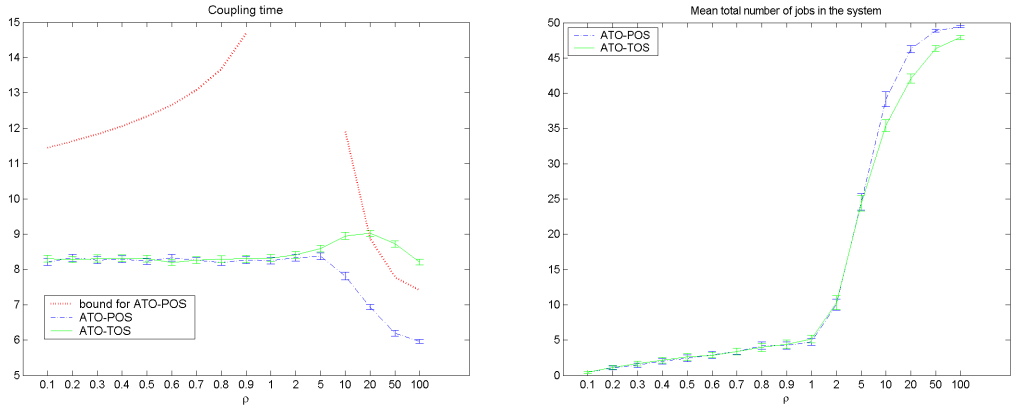


Figure 3 On the left: Coupling times for ATO-POS and ATO-TOS, and the upper bound for ATO-POS (we display $\log_2(T)$ where T is the mean coupling time), together with the 95% confidence intervals. On the right: Mean number of jobs in the system for ATO-POS and ATO-TOS, together with the 95% confidence intervals.

In Section 6, we will consider ATO systems with joint services. We explain in 6.1 why the existing methods (PSA or EPSA) cannot be used to sample (with reasonable complexity) the steady state distribution. Hence we introduce a new method, that we explain in the following section.

5 Aggregated envelopes

We introduce a new method of perfect sampling, that we apply in Section 6 to an Assemble-To-Order system with Partial Order Service (ATO-POS) and joint returns of items, also called the service tools model by [Vliegen and van Houtum, 2009]. In this model, the customers demand (or borrow) subsets of items and return them together (see Section 6.1 for more details). Although we are usually interested only in the total number of available items of each type, this information is not sufficient to describe the evolution of the system: In order to get a Markov chain, we need to keep track of the way items leave the stock, as they will be returned together. The state space becomes rapidly intractable: Even the dimension of the state space is exponential with respect to the number I of different item types. Thus even storing the vector representing the state of the system becomes challenging.

The idea of aggregated envelope method is to consider the projection of the state space on a more tractable space: In ATO-POS system with joint returns, we consider the projected space $\mathcal{X}$ of vectors $x = (x_1, \dots, x_I)$, where x_i is the total number of available type i items. Considering space $\mathcal{X}$ reduces exponentially the dimension of the state space. However, vector x does not contain all the information about the evolution of the system so we will need to construct a bounding chain that takes into account all possible evolutions.

More generally, we assume that our initial Markov chain is given by a tuple $\mathcal{M} = (\mathcal{N}, \mathcal{E}, \nu, g)$. We assume further that there is a projection function $\psi : \mathcal{N} \to \mathcal{X}$ such that $(\mathcal{X}, \preceq)$ is a finite lattice. In practice, $|\mathcal{X}|$ will be much smaller that $|\mathcal{N}|$. The state space $\mathcal{N}$ is not necessarily a lattice (we do not assume any ordering relation on $\mathcal{N}$). This is another important motivation for the aggregated envelope method: The state space of the ATO-POS system with joint returns is not a lattice for the product order.

In this section, we develop a method that samples an interval of $\mathcal{X}$ containing the projection of a state distributed according to the stationary distribution π of the original Markov chain. We will see that, under some conditions (see Section 6.7), it is even possible to sample a state in $\mathcal{N}$, distributed according to the stationary distribution.

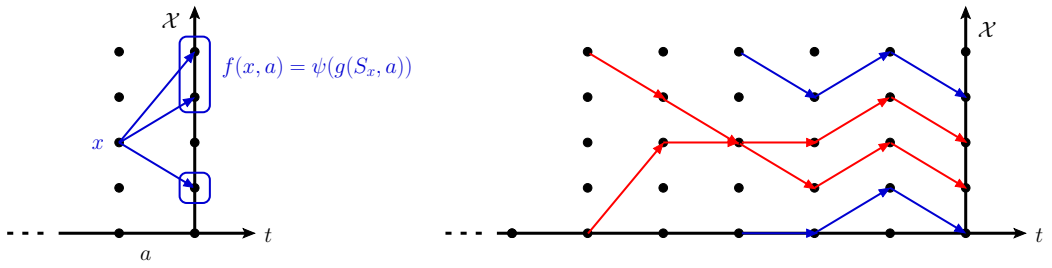


Figure 4 On the right: The idea of aggregation (we consider the evolution from *all* the states $n \in \mathcal{N}$ such that $\psi(n) = x$). On the left: Even if at some time $-t$ we have only one value for the projected process, this is not necessarily the case for times $-s$, $-t \leq -s \leq 0$, as this single projected value at time $-t$ can correspond to many different states in $\mathcal{N}$.

5.1 Aggregation

Our starting idea is to use the projected state space $\mathcal{X}$ for simulations. Intuitively, the original Markov chain $X = (X_t)_{t\in\mathbb{N}}$ evolves in $\mathcal{N}$, but we can only observe its projection $Y = (Y_t)_{t\in\mathbb{N}} = (\psi(X_t))_{t\in\mathbb{N}}$. Assume that the original chain is in state $n_0 \in \mathcal{N}$. The only information we have is $x = \psi(n_0) \in \mathcal{X}$. When an event $a \in \mathcal{E}$ occurs, we need to determine the next state in $\mathcal{X}$. As Y is not a Markov chain, we cannot determine the next state only from knowing x. Instead, we will consider the evolution from *all* the states $n \in \mathcal{N}$ such that $\psi(n) = x$. More formally, for $x \in \mathcal{X}$, we consider the following *subset* $S_x \subset \mathcal{N}$:

$$S_x = \{n \in \mathcal{N}, \psi(n) = x\} = \psi^{-1}(\{x\}).$$

and define the following function $f : \mathcal{X} \times \mathcal{E} \to \mathcal{P}(\mathcal{X})$, as illustrated in Figure 4:

$$f(x,a) = \psi\left(g(S_x,a)\right), \tag{9}$$

where g is the transition function of the original Markov chain.

Let $\mathcal{P}(\mathcal{X})$ denote the family of subsets of $\mathcal{X}$, and $F : \mathcal{P}(\mathcal{X}) \times \mathcal{E} \to \mathcal{P}(\mathcal{X})$ be the transition function defined by:

$$F(U,a) = \cup_{x\in U} f(x,a). \tag{10}$$

Since F is a bounding chain for the projected process, we could use a coupling from the past scheme (for one trajectory starting from $U = \mathcal{X}$) to provide an interval containing the projection of a state distributed according to the stationary distribution π: This is Algorithm 2. However, as explained below, this algorithm has several drawbacks and we will rather combine it with the method of envelopes developed in [Bušić et al., 2008].

The lack of knowledge induced by the projection on $\mathcal{X}$ forces us to consider *all* the states $n \in \mathcal{N}$ with the same projection x. This induces two main problems:

- Even if the original system couples, we may never have $|U| = 1$ in Algorithm 2.
- Even if at some time $-t$ we have only one value for the projected process (i.e. $|U| = 1$), this is not necessarily the case for times $-s$, $-t \leq -s \leq 0$, as this single projected value at time $-t$ can correspond to many different states in $\mathcal{N}$, as illustrated in Figure 4.

The first problem is similar as the case of EPSA algorithm when the envelope chain does not couple (i.e. the assumptions of Theorem 3 do not hold). The second problem could not occur for EPSA algorithm (a singleton interval of an envelope chain contains only a single possible state of the initial chain, so only one possible trajectory from that time until time 0). However, the two problems are of similar nature and the approaches described in Remark 4 remain valid. We will be interested in particular in performance bounds for increasing cost functions, described in Section 2.3.

Algorithm 2: A bounding chain based on aggregation

Data: *I.i.d.* events $(a_{-t})_{t\in\mathbb{N}} \in \mathcal{E}^{\mathbb{N}}$.
Result: Subset $U \subset \mathcal{X}$ containing the projection of a state $n^* \in \mathcal{N}$ distributed according to π.

```
begin
    t = 1; c = 0;
    repeat
        U = X;
        for i = t − 1 downto 0 do
            U := F(U, a_{−i}) ;
            if |U| = 1 then c = 1;
        t := 2t;
    until c = 1;
    return U;
```

In addition to these problems, space $\mathcal{X}$ can be too large to consider all the initial states $x \in \mathcal{X}$. Our objective is twofold. First, we would like to be able to compute $\psi\left(g(S_x, a)\right)$, for a given state x and event a, without considering all $n \in S_x$ (for ATO systems with joint services, $|S_x|$ grows exponentially with the number of item types). Second, we do not want to be forced to calculate, at each step, $\psi\left(g(S_x, a)\right)$ for *all* the current states x (even if we are able to calculate it easily for each state). To overcome this, we will combine the idea of aggregation with the method of envelopes developed in [Bušić et al., 2008].

5.2 Combining aggregation and envelopes

We will define a Markov chain $Y^{\inf}$ (respectively $Y^{\sup}$) that maps x to the infimum (resp. supremum) of $f(x, a) = \psi\left(g(S_x, a)\right)$. Consider the following transition functions: for all $x \in \mathcal{X}$ and $a \in \mathcal{E}$,

$$g^{\inf}(x, a) \stackrel{\text{def}}{=} \inf f(x, a), \quad g^{\sup}(x, a) \stackrel{\text{def}}{=} \sup f(x, a).$$

Let $m, M \in \mathcal{X}$ such that $m \preceq M$. The envelope method in [Bušić et al., 2012a] changes the subset $[m, M]$ into a new subset $[m', M']$ (that depends on the event $a \in \mathcal{E}$ and that usually involves only the transition function of one Markov chain). Here we consider two Markov chains $Y^{\inf}$ and $Y^{\sup}$ on the same space $\mathcal{X}$, with the same set of events $\mathcal{E}$, but with two different transition functions $g^{\inf}$ and $g^{\sup}$. Considering separately the envelopes of the infimum and the supremum chains does not necessarily sandwich the projected process (if either $g^{\inf}$ or $g^{\sup}$ is not monotone). We define the aggregated envelope transition function as follows: For $m, M \in \mathcal{X}$ such that $m \preceq M$, and $a \in \mathcal{E}$,

$$\begin{aligned} H\left([m, M], a\right) &\stackrel{\text{def}}{=} \left[\inf_{m \preceq x \preceq M} g^{\inf}(x, a), \sup_{m \preceq x \preceq M} g^{\sup}(x, a)\right] \\ &= \left[\underline{H}\left([m, M], a\right), \overline{H}\left([m, M], a\right)\right]. \end{aligned}$$

In order to compare the projected process $Y = \psi(X)$ of the original chain X to the lower and the upper envelopes of H, we need the following notation. Assume the sequence of events $a_{-t+1}, \ldots, a_0 \in \mathcal{E}$ fixed. For $(t, n) \in \mathbb{N} \times \mathcal{N}$, $X(-t, n)$ stands for a realization of X that starts from n at time $-t$, while $X_{-s}(-t, n)$ denotes the value of this realization at time $-s$. The next proposition shows that the chain with transition function H is a bounding interval chain for the projected process $\psi(X)$.

Proposition 16. *Let $n \in \mathcal{N}$ and $y, z \in \mathcal{X}$ such that $y \preceq \psi(n) \preceq z$. Let $t \in \mathbb{N}$ and $a_{-t+1}, \ldots, a_0 \in \mathcal{E}$. Then we have for any $s \leq t$:*

$$\begin{aligned} \underline{H}\left([y,z], a_{-t+1\to -s}\right) &\preceq \psi(X_{-s}(-t,n)) \\ &\preceq \overline{H}\left([y,z], a_{-t+1\to -s}\right). \end{aligned}$$

Proof. Let $x := \psi(n)$. We prove the result by descending induction on s. First, for $s = t$ the result is trivial, as $a_{-t+1\to -s} = \epsilon$ (empty word), so

$$\underline{H}\left([y,z], \epsilon\right) = y \preceq \psi(X_{-t}(-t,n)) \preceq z = \overline{H}\left([y,z], \epsilon\right).$$

Assume the result is true for some s, $s \leq t$. Let $n' := X_{-s}(-t,n)$, $x' := \psi(n')$, $y' := \underline{H}\left([y,z], a_{-t+1\to -s}\right)$ and $z' := \overline{H}\left([y,z], a_{-t+1\to -s}\right)$. Then we have that $y' \preceq x' \preceq z'$ by induction hypothesis. For $s-1$, we have:

$$\begin{aligned} \underline{H}\left([y',z'], a_{-s+1}\right) &\preceq g^{\inf}(x', a_{-s+1}) \preceq \psi(X_{-s+1}(-t,n)) \\ &\preceq g^{\sup}(x', a_{-s+1}) \preceq \overline{H}\left([y',z'], a_{-s+1}\right). \end{aligned}$$

By the definition of y' and z', $\underline{H}\left([y',z'], a_{-s+1}\right) = \underline{H}\left([y,z], a_{-t+1\to -s+1}\right)$ and $\overline{H}\left([y',z'], a_{-s+1}\right) = \overline{H}\left([y,z], a_{-t+1\to -s+1}\right)$, which gives the result for $s-1$ and ends the proof. □

The aggregated envelope method is summarized in Algorithm 3.

Algorithm 3: Aggregated Envelope Perfect Sampling

Data: *I.i.d.* events $(a_{-t})_{t\in\mathbb{N}} \in \mathcal{E}^{\mathbb{N}}$.
Result: Interval $[m^*, M^*] \subset \mathcal{X}$ containing the projection of a state $n^* \in \mathcal{N}$ distributed according to π.

```
begin
    t = 1; c = 0;
    repeat
        m := ⊥ (∈ X); M := ⊤ (∈ X);
        for i = t − 1 downto 0 do
            [m, M] := H([m, M], a_{−i}) ;
            if m = M then c = 1;
        t := 2t;
    until c = 1;
    m* := m; M* := M;
    return m*, M*;
```

5.3 Alternative stopping condition

Note that Algorithm 3 only gives an interval that contains the projection of a state distributed according to the stationary distribution. In Algorithm 4, we use a heuristic to relax the stopping condition: Instead of stopping when the upper and lower envelopes meet ($m = M$), we stop when they meet at least once on each component between time $-t+1$ and time 0. The intuition behind this is the following: For a chain which is componentwise coupling (see Definition 7 in Section 3.2), this condition is sufficient to insure that the chain has coupled. In the general case, this algorithm can be faster than Algorithm 3 (and still obtain bounds on the projection of a steady state). In addition, it is easier to bound the stopping time of Algorithm 4. However, as the chain has not necessarily coupled, there can be cases for which bounds obtained using Algorithm 4 are loose.

In the case of the ATO-POS system with joint returns, we are able to provide some bounds for the stopping time of Algorithm 4 (see Theorem 22 in Section 6.6). Moreover, we will see that the results obtained for Algorithms 3 and 4 are very close.

Algorithm 4: Modified stopping condition

Data: *I.i.d.* events $(a_{-t})_{t\in\mathbb{N}} \in \mathcal{E}^{\mathbb{N}}$.
Result: Interval $[m^*, M^*] \subset \mathcal{X}$ containing the projection of a state $n^* \in \mathcal{N}$ distributed according to to the stationary distribution.

```
begin
    t := 1; c := zeros(1, I);
    repeat
        m := ⊥ (∈ X); M := ⊤ (∈ X);
        for i = t − 1 downto 0 do
            [m, M] := H([m, M], a_{−i}) ;
            for j = 1 to I do
                if m(j) = M(j) then c(j) := 1;
        t := 2t;
    until c = ones(1, I);
    m* := m; M* := M;
    return m*, M*;
```

6 ATO systems with joint returns

The case of ATO systems with joint returns is considerably more difficult, due to the need to track the information on which items were sent together to the customer. This leads to extremely large state space, not having a natural lattice structure (as will be shown in Section 6.1), which makes it intractable for PSA or EPSA algorithms. We show in this section how to use the aggregated envelope method proposed in Section 5 to reduce the state space and overcome the lack of lattice structure.

Even if the general idea is also valid for ATO-TOS, we focus only on ATO-POS for two reasons: unlike the TOS case, POS model does not have a product form solution, which makes it even more challenging. The second reason is that the monotonicity of arrivals in the POS case makes computations easier to some extent and simplifies the presentation of the general ideas. The adaptation for the TOS case is discussed in Section 7 (although TOS has product form, the estimation of the normalizing constant is still a hard problem).

We start by defining the model and its event representation in Sections 6.1 and 6.2. In order to compute the aggregated envelope chain H (Section 6.5), we first analyze separately the supremum chain (Section 6.3) and the infimum chain (Section 6.4). In Section 6.6, we give a bound on the mean stopping time of Algorithm 4. When the service rate is high (Section 6.7), one can obtain perfect samples from the steady state distribution of the Markov chain, with an algorithm whose complexity is quadratic with respect to the total capacity $|C| = \sum_{i=1}^{I} C_i$.

6.1 Model description

We consider an Assemble-To-Order system with Partial Order Service (ATO-POS) and joint returns of items, also called the service tools model by [Vliegen and van Houtum, 2009]. As in Section 3, demands for each subset $A \subset \mathcal{I} = \{1, ..., I\}$ follow a Poisson process of rate λ_A. As before, we assume that the number $|\{A \subset \mathcal{I} : \lambda_A \neq 0\}|$ of subsets a customer can ask for is small (for instance linear with respect to the number I of item types). If some demanded items are not available, then the customer takes the available items (POS case). The available items return from the customer together (unlike the individual replenishment assumption in previous sections), and after an exponential time of rate μ.

In terms of a network of I queues, this means that we consider joint services: items that arrived in the system together (taken by one customer) will also leave the system together (returned from the customer). Therefore we need to keep track of which items arrived to the queues together.

The system can be modeled as a continuous time Markov chain with state space:

$$\mathcal{N} = \left\{ (n_A)_{A\subseteq\mathcal{I}, A\neq\emptyset} \,:\, \forall A, n_A \geq 0 \;\&\; \forall i, \sum_{A: i\in A} n_A \leq C_i \right\}.$$

where n_A is the number of subsets A currently at the customers. Let e_A denote the vector of $\mathcal{N}$ whose coordinate A is equal to 1, and others are 0 (this notation also stands for $\sum_{i\in A} e_i \in \mathcal{X}$, but no confusion will be possible since the space will always be specified). We define the projection ψ on space $\mathcal{X} = \{0, ..., C_1\} \times \{0, ..., C_2\} \times \cdots \times \{0, ..., C_I\}$:

$$\psi : \left\{ \begin{array}{ccc} \mathcal{N} & \longrightarrow & \mathcal{X} \\ n = (n_A)_{A\subset\mathcal{I}} & \longmapsto & x = \left(\sum_{A: i\in A} n_A\right)_{i\in\mathcal{I}} \end{array} \right.$$

The total number of items of each type in queues (i.e. currently used by the customers) is given by a vector $x = (x_1, \ldots, x_I) \in \mathcal{X}$, where x_i is the number of items of type i. We will consider the product order $\leq$ on $\mathcal{X}$.

Note that the only dimension of the state space $\mathcal{N}$ is $2^I - 1$. In addition, it is easy to see that $\mathcal{N}$ equipped with the usual product order is not a lattice, as the supremum of two points can exceed the state space. For instance, let $I = 2$ and set $c := \min(C_1, C_2)$. Let $n, m \in \mathcal{N}$ such that $n_{\{1,2\}} = c$, $n_{\{1\}} = n_{\{2\}} = 0$ and $m_{\{1\}} = m_{\{2\}} = c$, $m_{\{1,2\}} = 0$. Then $\psi(n) = \psi(m) = (c, c)$, yet the supremum $M = \sup(n, m)$ satisfies: $\psi(M) = (2c, 2c)$. Therefore, PSA or EPSA cannot be applied directly using the product order on $\mathcal{N}$: Even with monotone events, it is impossible to upper-bound any given subset of states by only a few extremal states.

We have two different types of transitions. For each $n \in \mathcal{N}$, and for each $A \subset \mathcal{I}$:

- There is a demand for subset A, with rate λ_A. The new state is: $n + e_{A^{(n)}}$, where

$$A^{(n)} = \{i \in A : (\psi(n))_i < C_i\} \tag{11}$$

 denotes the items of set A that are available in state n and that are sent together to the customer.

- If $n_A > 0$, there is a joint service of A, with rate $\mu \cdot n_A$. The new state is $n - e_A$.

By a standard uniformization procedure, we can transform the above continuous time Markov chain to a discrete time Markov chain. The outgoing rate for each state is upper-bounded by $\Lambda := \lambda + \mu|C|$. We take the uniformization constant to be equal to Λ. We now give a discrete event representation of the uniformized Markov chain.

6.2 Discrete event representation

Arrivals. For any $A \subset \mathcal{I}$, $A \neq \emptyset$, let d_A be the event of probability λ_A/Λ that corresponds to a "joint arrival to queues in A". We give the transition function g of the Markov chain X on $\mathcal{N}$ for an arrival d_A, $A \subset \mathcal{I}$, $A \neq \emptyset$, in state $n \in \mathcal{N}$:

$$g(n, d_A) = n + e_{A^{(n)}}, \tag{12}$$

where $A^{(n)} \subset A$ is defined in (11).

Services. Unfortunately, if service events are not well chosen, the supremum chain does not move with any service. Indeed, let us observe a fixed $i \in \mathcal{I}$. In order to have $(g^{\sup}(x, a))_i = x_i - 1$, for some service event a, in *all* states $n \in \mathcal{N}$ such that $\psi(n) = x$, event a must correspond to a service of some subset $A \subset \mathcal{I}$ that contains i (A can depend on n). Otherwise, if there is at least one state n whose i-th queue is not served, then the i-th component of the supremum does not move. This makes the definition of services a little tricky. In addition, the following representation is linear with respect to $|C|$.

Before defining service events, we need to define an ordering for the non-empty subsets of $\mathcal{I}$. For all $i \in \mathcal{I}$, we define $\left(A_k^i\right)_{0\leq k\leq 2^{I-i}-1}$ as an ordering of all the subsets of $\{i, ..., I\}$ containing i (subsets $\left(A_k^1\right)_k$ are those that contain 1, $\left(A_k^2\right)_k$ those that contain 2 but not 1, and so on). More precisely, let $i \in \mathcal{I}$ and $k \in \{0, 1, ..., 2^{I-i} - 1\}$. We set $k = k_1 ... k_{I-i}$ for the binary representation of k ($k = \sum_{s=1}^{I-i} k_s 2^{I-i-s}$, where $k_s \in \{0, 1\}$). Then the subset A_k^i is by definition such that:

- $i \in A_k^i$;
- for $s \in \{i + 1, ..., I\}$, $s \in A_k^i$ if and only if $k_{s-i} = 0$.

For instance, if $I = 4$, $A_0^1 = \{1, 2, 3, 4\}$, $A_1^1 = \{1, 2, 3\}$, ..., $A_7^1 = \{1\}$, $A_0^2 = \{2, 3, 4\}$, ..., $A_0^4 = \{4\}$.

Now we can define services. Let r_j^i, $1 \leq i \leq I$, $1 \leq j \leq C_i$, be independent events of probability μ/Λ, such that for $n \in \mathcal{N}$:

- If $\sum_{\ell=0}^{2^{I-i}-1} n_{A_\ell^i} < j$, set $g(n, r_j^i) = n$;
- If $\sum_{\ell=0}^{2^{I-i}-1} n_{A_\ell^i} \geq j$, let k be the minimal element of $\{0, 1, ..., 2^{I-i}-1\}$ such that $\sum_{\ell=0}^{k-1} n_{A_\ell^i} < j \leq \sum_{\ell=0}^{k} n_{A_\ell^i}$. Set $g(n, r_j^i) = n - e_{A_k^i}$.

Note that the total number of events corresponding to services is exactly $|C|$. The next lemma shows that this definition of services agrees with the rates of the Markov chain defined in 6.1.

Lemma 17. *For any state $n \in \mathcal{N}$ and any subset A_k^i, $1 \leq i \leq I$, $0 \leq k \leq 2^{I-i} - 1$, with the above definition of services, the total probability to go from state n to state $n - e_{A_k^i}$ is $\frac{\mu}{\Lambda} \cdot n_{A_k^i}$.*

Proof. Let $n \in \mathcal{N}$ and $i \in \mathcal{I}$ be fixed. We assume that $n \neq 0$ (otherwise the result is obvious). First we explain the effect of services $(r_j^i)_j$ on n. For any $j \in \{1, ..., C_i\}$, the service r_j^i corresponds to a joint service of A_k^i, for some k that depends on state n. More precisely, for $j = 1$, service r_1^i corresponds to a joint service of the subset A_k^i, where k is the smallest integer such that $n_{A_k^i} > 0$. We set $n' = n - e_{A_k^i}$. If $n' = 0$, service r_1^i is the only service that modifies n. Otherwise, for $j = 2$, service r_2^i corresponds to a joint service of the subset A_ℓ^i, where ℓ is the smallest integer such that $n'_{A_\ell^i} = n_{A_\ell^i} - \mathbb{1}_{\{\ell=k\}} > 0$. Set $n'' = n' - e_{A_\ell^i}$. If $n'' = 0$, then r_1^i and r_2^i are the only services that change n, and if $n'' \neq 0$, we continue as before.

The probability to go from state n to $n - e_{A_k^i}$, for any $0 \leq k \leq 2^{I-i} - 1$ is:

$$\sum_{j=1}^{C_i} \mathbb{P}\left(g(n, r_j^i) = n - e_{A_k^i}\right) = \frac{\mu}{\Lambda} \left| \left\{ j \in \{1, \ldots, C_i\} : \sum_{\ell=0}^{k-1} n_{A_\ell^i} < j \leq \sum_{\ell=0}^{k} n_{A_\ell^i} \right\} \right| = \frac{n_{A_k^i} \mu}{\Lambda}.$$

□

We want to compute the aggregated envelope transition function H (defined in 5.2). For this, we first need to compute the supremum chain $g^{\sup}$ and the infimum chain $g^{\inf}$ (defined in 5.2). We begin by the supremum chain.

6.3 Supremum chain

The supremum chain was defined as follows in Section 5.2. For all $a \in \mathcal{E}$:

$$g^{\sup}(x, a) = \sup \{\psi(g(n, a)) : n \in \mathcal{S},\ \psi(n) = x\}.$$

Lemma 18. *For all $x \in \mathcal{X}$, set $\hat{x}_i \stackrel{\text{def}}{=} \max\{x_i - (x_1 + ... + x_{i-1}), 0\}$. Then the transition function of the supremum chain is, for $A \subset \mathcal{I}$, $A \neq \emptyset$, $i \in \mathcal{I}$, $1 \leq j \leq C_i$, and $x \in \mathcal{X}$:*

$$\begin{cases} g^{\sup}(x, d_A) &= x + \sum_{k \in A} \mathbb{1}_{\{x_k < C_k\}} e_k = f_P(x, d_A), \\ g^{\sup}(x, r_j^i) &= x - \mathbb{1}_{j \leq \hat{x}_i} e_i. \end{cases}$$

where f_P is the transition function for the model ATO-POS with individual returns.

Proof. Arrivals. Let $A \subset \mathcal{I}$, $A \neq \emptyset$, $x \in \mathcal{X}$, and $n \in \mathcal{N}$ such that $\psi(n) = x$. Applying ψ to (12) gives that: $\psi(g(n, d_A)) = x + \sum_{i \in A} \mathbb{1}_{\{x_i < C_i\}} e_i$, hence the set $f(\{x\}, d_A) = \left\{x + \sum_{i \in A} \mathbb{1}_{\{x_i < C_i\}} e_i\right\}$ has only one element (see (10) for the definition of f), and $g^{\sup}(x, d_A) = x + \sum_{i \in A} \mathbb{1}_{\{x_i < C_i\}} e_i = g^{\inf}(x, d_A)$.

Services. Let $i \in \mathcal{I}$, $1 \leq j \leq C_i$, and $x \in \mathcal{X}$. We first assume $j \leq \hat{x}_i$ (note that this implies $x_i \geq 1$). For n such that $\psi(n) = x$, we will prove that

$$\left(\psi(g(n, r_j^i))\right)_i = x_i - 1, \tag{13}$$

which implies that $g^{\sup}(x, r_j^i) \leq x - e_i$. In addition, if $n_{\{i\}} = x_i$, then $\psi\left(g(n, r_j^i)\right) = x - e_i$. So $g^{\sup}(x, r_j^i) = x - e_i$. Now we prove (13). Let $n \in \mathcal{N}$ such that $\psi(n) = x$. Then

$$x_i = \sum_{A: i \in A} n_A = \sum_{i'=1}^{i} \sum_{k: i \in A_k^{i'}} n_{A_k^{i'}} \leq x_1 + ... + x_{i-1} + \sum_k n_{A_k^i}.$$

where the second equality comes from the recursive definition of subsets $\left(A_k^i\right)_k$. Hence we have: $\sum_k n_{A_k^i} \geq \max\{x_i - (x_1 + ... + x_{i-1}), 0\} = \hat{x}_i \geq j$, so we get the result by the definition of the event r_j^i.

We prove the second case now: assume $j > \hat{x}_i$. For n such that $n_{\{k\}} = x_k$ for all $1 \leq k \leq I$ (and $n_A = 0$ if $|A| \geq 2$), we have: $g(n, r_j^i) = n$, so $g^{\sup}(x, r_j^i) = x$, which concludes the proof.

Monotonicity.

As a corollary of Lemma 18, we have the following result (the proof for services follows from the fact that $x \leq y$ and $x_i = y_i$ imply $\hat{x}_i \geq \hat{y}_i$).

Proposition 19. *Let $a \in \mathcal{E}$ be any arrival or service. Under the usual product order on $\mathcal{X}$, the event a is monotone for the supremum chain in the ATO-POS system with joint services, i.e. for all $x, y \in \mathcal{X}$ such that $x \leq y$, we have that: $g^{\sup}(x, a) \leq g^{\sup}(y, a)$.*

6.4 Infimum chain

The infimum chain was defined as follows in Section 5.2, for all $x \in \mathcal{X}$, $a \in \mathcal{E}$:

$$g^{\inf}(x, a) = \inf\left\{\psi\left(g(n, a)\right) : n \in \mathcal{N},\ \psi(n) = x\right\}.$$

For arrivals, we have for any $A \subset \mathcal{I}$, $A \neq \emptyset$, and $x \in \mathcal{X}$ (see the proof of Lemma 18):

$$g^{\inf}(x, d_A) = x + \sum_{i \in A} \mathbb{1}_{\{x_i < C_i\}} e_i = g^{\sup}(x, d_A).$$

The following lemma describes the infimum chain for services.

Lemma 20. *Let $p \in \mathcal{I}$. The p-th component of $g^{\inf}(x, r_j^i)$ satisfies:*

- *If $p < i$, then $\left(g^{\inf}(x, r_j^i)\right)_p = x_p$;*
- *If $p = i$, then $\left(g^{\inf}(x, r_j^i)\right)_p = x_i - \mathbb{1}\{j \leq x_i\}$;*
- *If $p > i$, then $\left(g^{\inf}(x, r_j^i)\right)_p = x_p - \mathbb{1}\left\{x_p > 0\ \&\ j \leq \min\left(\sum_{i'=i+1}^{p} x_{i'}, x_i\right)\right\}$.*

Proof. Since we consider the product order on $\mathcal{X}$, the infimum can be computed componentwise, *i.e.* for $p \in \mathcal{I}$: $\left(g^{\inf}(x, r_j^i)\right)_p = \inf\left\{\left(\psi\left(g(n, r_j^i)\right)\right)_p : n \in \mathcal{N},\ \psi(n) = x\right\}$.

Case $p < i$. By the definition of r_j^i, there is no change on the p-th component, so $\left(g^{\inf}(x, r_j^i)\right)_p = x_p$.

Case $p = i$. If $j > x_i$, then no n such that $(\psi(n))_i = x_i$ can be modified by r_j^i, so $\left(g^{\inf}(x, r_j^i)\right)_i = x_i$. If $j \leq x_i$, let us choose n such that $\psi(n) = x$ and $n_{\{i\}} = x_i$. Then $\left(g(n, r_j^i)\right)_i = x_i - 1$, so $\left(g^{\inf}(x, r_j^i)\right)_i = x_i - 1$.

Case $p > i$. Clearly, a state $n \in \psi^{-1}(\{x\})$ can be modified by r_j^i only if $x_p > 0$. Moreover, its p-th component is modified by r_j^i if and only if there exists $k \in \{0, 1, ..., 2^{I-i} - 1\}$ such that

$$p \in A_k^i \quad \text{and} \quad \sum_{\ell=0}^{k-1} n_{A_\ell^i} < j \leq \sum_{\ell=0}^{k} n_{A_\ell^i}. \tag{14}$$

In particular, the condition $j \leq \sum_{\ell=0}^{2^{I-i}-1} n_{A_\ell^i}$ (and so $j \leq x_i$) is necessary for n to change.

Assume $j > \sum_{i'=i+1}^{p} x_{i'}$. We will show that no n such that $\psi(n) = x$ can change on the p-th component by r_j^i. Indeed, due to the recursive definition of subsets $(A_k^i)_k$, we have that A_k^i cannot contain p if $k \geq 2^{I-i-1} - 2^{I-i+1-p}$, and, for $k = 2^{I-i-1} - 2^{I-i+1-p} - 1$, we have:

$$\sum_{\ell=0}^{k} n_{A_\ell^i} = \sum_{i'=i+1}^{p} \sum_{\substack{\ell : i' \in A_\ell^i \\ 1 < i'' < i' \Rightarrow i'' \notin A_\ell^i}} n_{A_\ell^i} \leq \sum_{i'=i+1}^{p} x_{i'} < j.$$

So (14) cannot hold, and n cannot change on the p-th component.

Assume now $x_p > 0$ and $j \leq \min\left(\sum_{i'=i+1}^{p} x_{i'}, x_i\right)$. We give the construction of $n \in \mathcal{N}$ such that $\psi(n) = x$ and $\left(g(n, r_j^i)\right)_p = x_p - 1$:

Construction of n
$m := \min(x_i, x_p)$;
$V := 0$; $a := 0$;
for $i' = i + 1$ **to** $p - 1$ **do**
$\quad a := \min(x_{i'}, j - 1 - V, x_i - m - V)$;
$\quad n_{\{i,i'\}} := a$; $\quad n_{\{i'\}} := x_{i'} - a$;
$\quad V := V + a$;
end
$n_{\{i,p\}} := m$; $\quad n_{\{p\}} := x_p - m$; $\quad n_{\{i\}} := x_i - V - m$;
$n_{\{i'\}} = x_{i'}$ for $i' < i$ or $i' > p$, and $n_A = 0$ for other subsets A ;

Clearly, $\psi(n) = x$. Note also that our hypotheses imply that $m \geq 1$. Let $k \in \{0, 1, ..., 2^{I-i} - 1\}$ such that $A_k^i = \{i, p\}$. We will show that: $\sum_{\ell=0}^{k-1} n_{A_\ell^i} < j \leq \sum_{\ell=0}^{k} n_{A_\ell^i}$, which will end the whole proof (cf. (14)). Let $\tilde{V}$ be the value of V at the end of the loop. Then we have that $\sum_{\ell=0}^{k-1} n_{A_\ell^i} = \tilde{V}$ and $\sum_{\ell=0}^{k} n_{A_\ell^i} = \tilde{V} + m$, so we are left to prove that:

$$\tilde{V} < j \leq \tilde{V} + m \tag{15}$$

We study more precisely what happens during the loop. If there exists a step during the loop such that $a = j - 1 - V$ (resp. $a = x_i - m - V$), then $a = 0$ in all the next steps and, at the end of the loop, $\tilde{V} = j - 1$ (resp. $\tilde{V} = x_i - m$). In these two cases, (15) follows easily. The remaining case is the one such that, at each step of the loop, $a = x_{i'} \leq \min(j - 1 - V, x_i - m - V)$. In particular, considering the last step of the loop, the previous inequality gives that $\tilde{V} \leq j - 1$. Moreover, $\tilde{V} = \sum_{i'=i+1}^{p-1} x_{i'}$ in this case, so the fact that $j \leq \min\left(\sum_{i'=i+1}^{p} x_{i'}, x_i\right)$ gives that $j \leq \tilde{V} + m$ (considering the two possible cases for the value of m). □

Non-monotonicity.

Unfortunately, services are not monotone for the infimum chain. Indeed, let $I = 2$, $x = (0, 1)$, and $y = (1, 1)$. Then $x \leq y$, yet $g^{\inf}(x, r_1^1) = (0, 1) \geq (0, 0) = g^{\inf}(y, r_1^1)$.

6.5 The aggregated envelope chain H

We use Sections 6.3 and 6.4 to compute the aggregated envelope chain H (defined in Section 5.2). Let $m, M \in \mathcal{X}$, $m \leq M$. Our goal is to compute, for all $a \in \mathcal{E}$:

$$\begin{aligned} H\left([m,M],a\right) &= \left[\inf_{m\leq x\leq M} g^{\inf}(x,a), \sup_{m\leq x\leq M} g^{\sup}(x,a)\right] \\ &= \left[\underline{H}\left([m,M],a\right), \overline{H}\left([m,M],a\right)\right]. \end{aligned}$$

Arrivals. Let $A \subset \mathcal{I}$, $A \neq \emptyset$. Then $H\left([m,M],d_A\right) = [f_P(m,d_A), f_P(M,d_A)]$ since $g^{\inf}(x,d_A) = g^{\sup}(x,d_A) = x + \sum_{k\in A} \mathbb{1}_{\{x_k<C_k\}} e_k = f_P(x,d_A)$ for all $x \in \mathcal{X}$. In addition, f_P is monotone due to Proposition 6, which gives the result for H.

Services. Let $i \in \mathcal{I}$ and $j \in \{1,2,...,C_i\}$. By Proposition 19, $g^{\sup}$ is monotone, thus $\overline{H}\left([m,M],r_j^i\right) = g^{\sup}(M,r_j^i) = M - \mathbb{1}_{j\leq \hat{M}_i} e_i$. In order to compute $\underline{H}\left([m,M],r_j^i\right)$, we need to compute the lower envelope of the infimum chain (Lemma 21).

Lemma 21. *Let m, $M \in \mathcal{X}$ such that $m \leq M$. Let $i \in \mathcal{I}$ and $j \in \{1,2,...,C_i\}$. Set $m' = \underline{H}\left([m,M],r_j^i\right)$. We compute each component p of m', for $p \in \mathcal{I}$, and we can distinguish three cases:*

- *If $p < i$, then $m'_p = m_p$;*
- *If $p = i$, then $m'_p = m_i - \mathbb{1}\{j \leq m_i\}$;*
- *If $p > i$, then $m'_p = m_p - \mathbb{1}\left\{m_p > 0 \ \&\ j \leq \min\left(\sum_{i'=i+1}^{p-1} M_{i'} + m_p, M_i\right)\right\}$.*

Proof. Case $p < i$. By the definition of r_j^i, there is no change on the p-th component, so $m'_p = m_p$.

Case $p = i$. Let $x, y \in \mathcal{X}$ such that $x \leq y$. Then: $\left(g^{\inf}(x,r_j^i)\right)_p = x_i - \mathbb{1}\{j \leq x_i\} \leq y_i - \mathbb{1}\{j \leq y_i\} = \left(g^{\inf}(y,r_j^i)\right)_p$. Indeed, if $x_i < y_i$, this is obvious since y_i looses at most 1. Otherwise, $x_i = y_i$ and x and y move together.

Case $p > i$. If $j > M_i$, then for all $x \in \mathcal{X}$ such that $m \leq x \leq M$, we have that $j > x_i$, so x cannot move and $m'_p = m_p$. Assume $j \leq M_i$ from now on. If $m_p = 0$, then $m'_p = 0$, so we assume $m_p > 0$. We distinguish two cases:

- Assume $j > \sum_{i'=i+1}^{p-1} M_{i'} + m_p$. Let $x \in \mathcal{X}$ such that $m \leq x \leq M$. If $x_p > m_p$, then $\left(g^{\inf}(x,r_j^i)\right)_p \geq m_p$. Otherwise, $x_p = m_p$, and so $j > \sum_{i'=i+1}^{p-1} M_{i'} + m_p \geq \sum_{i'=i+1}^{p} x_{i'}$. Hence $\left(g^{\inf}(x,r_j^i)\right)_p = x_p$. Finally this leads to $m'_p = m_p$.

- Assume $j \leq \sum_{i'=i+1}^{p-1} M_{i'} + m_p$. We have that $m'_p \geq m_p - 1$, and we will define $x \in \mathcal{X}$ such that $m \leq x \leq M$ and $\left(g^{\inf}(x,r_j^i)\right)_p = m_p - 1$. This will then prove that $m'_p = m_p - 1$. Set $x_p := m_p$ and $x_{i'} := M_{i'}$ for all $i' \in \{1,2,...,I\} \setminus \{p\}$. Thus $j \leq \min\left(\sum_{i'=i+1}^{p} x_{i'}, x_i\right)$ and $\left(g^{\inf}(x,r_j^i)\right)_p = m_p - 1$.

□

Hence we gave the computation of the aggregated envelope chain H, which is necessary to use Algorithms 3 (in Section 5.2) and 4 (in Section 5.3).

In order to give a bound on the complexity of Algorithm 3 or 4, we have to take into account:

- the random choice of events,
- the computation of the aggregated envelope transition function,
- the stopping time of Algorithm 3 or 4.

The number of services is linear with respect to $|C|$. In addition, the number of arrivals is equal to the number of subsets A such that $\lambda_A \neq 0$: If this number is of order $O(I)$, using alias method by [**?**], the choice of events can be done in a constant time. In addition, Lemma 21 shows that the computation of the aggregated envelope transition function can be done in a linear time with respect to I. We next study the stopping time of Algorithm 4.

6.6 Bound for the stopping time of Algorithm 4

We suppose there exist two subsets $\mathcal{I}_0$ and $\mathcal{I}_C \subset \mathcal{I}$, $\mathcal{I} = \mathcal{I}_0 \cup \mathcal{I}_C$, such that:

(i) $\mu > \sum_{i\in\mathcal{I}_0} \lambda_i$,

(ii) $\delta_p \stackrel{\text{def}}{=} \lambda_p - \mu\left(\sum_{i=1}^{p} C_i - 1\right) > 0$ for all $p \in \mathcal{I}_C$.

Without loss of generality, we can change the numbering of queues such that:

(iii) $\Big(i \in \mathcal{I}_0 \text{ and } j \in \mathcal{I}_C \Big) \implies i \leq j$.

In the case where assumptions *(i)* and *(ii)* are not satisfied, we are not able to give a bound on the stopping time of Algorithm 4 (using the same method for the proof). The last condition comes from the expression of $\overline{H}\left([m,M], r_k^i\right) = M - \mathbb{1}_{k\leq \hat{M}_i} e_i$, where $\hat{M}_i = \max\{M_i - (M_1 + ... + M_{i-1}), 0\}$. Indeed we want the upper process to reach 0 on components $i \in \mathcal{I}_0$, and the lower process to reach C_j on components $j \in \mathcal{I}_C$. Yet the upper process cannot decrease on the i-th coordinate (due to service r_k^i, for any k) if there exists $j < i$ such that M_j is greater than M_i (in that case, $\hat{M}_i = 0$ and $\overline{H}\left([m,M], r_k^i\right) = M$).

Theorem 22. *Assume conditions (i) to (iii) hold. Then we can bound the time τ_{Alg} for which all components couple at least once by:*

$$\mathbb{E}\left[\tau_{Alg}\right] \leq \frac{\Lambda}{\mu - \sum_{i\in\mathcal{I}_0}\lambda_i} \sum_{i\in\mathcal{I}_0} C_i + \sum_{p\in\mathcal{I}_C} \frac{\Lambda}{\delta_p} C_p.$$

The proof is given in Appendix D, and requires Appendix C, in which we give a bound on the mean hitting to zero for the supremum chain $Y^{\sup}$.

We discuss the complexity of Algorithm 4 (with respect to $|C|$) in the case where $\mathcal{I}_0 = \mathcal{I}$ (high service rate case): Since $\Lambda = \lambda + \mu|C|$, we have that the mean stopping time of Algorithm 4 is of order $O\left(|C|^2\right)$ (using Theorem 22). Due to the discussion at the end of Subsection 6.5, the complexity of Algorithm 4 is thus quadratic with respect to $|C|$ (for high service rate).

Recall this is only a stopping criterion: The chain does not necessarily couple if we use Algorithm 4, and it is interesting to compare it with Algorithm 3, for which the stopping time is greater than or equal to the coupling time of the chain.

In Figure 5 on the left, we give stopping times for Algorithms 3 and 4 (ATO-POS with joint returns), for the following parameters: $I = 5$, $C_i = 10, \forall i$, $\lambda_A = \frac{1}{2^{|A|-1}}$, and $\mu_i(x_i) = \mu x_i$, with $\rho = \frac{\lambda_i}{\mu}$. The size of the sample is $N = 100$. We can observe that the mean stopping times of both algorithms are very close. In Figure 5 on the right, we provide mean distance between upper and lower bounding states at time 0 using 1-norm, *i.e.* $\sum_{i\in\mathcal{I}}(M_i^* - m_i^*)$.

6.7 Sampling the original (non aggregated) chain

Instead of having an interval that contains the projection of a state distributed according to the stationary distribution, we can even sample exactly the steady state distribution. However, the algorithm we provide here is efficient only in the high service rate case. We prove in Theorem 23 that its stopping time is quadratic with respect to the total capacity $|C|$.

The supremum chain $Y^{\sup}$ is monotone (Proposition 19, in Section 6.3). This gives directly the computation for the upper envelope of H: for any $m, M \in \mathcal{X}, a \in \mathcal{E}$, $\overline{H}\left([m,M], a\right) = g^{\sup}(M, a)$.

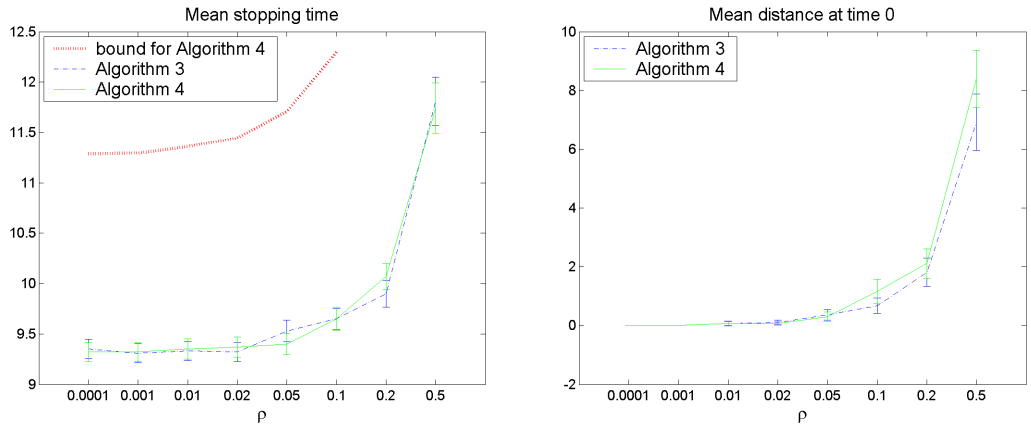


Figure 5 On the left: Stopping times for Algorithms 3 and 4 (ATO-POS with joint returns), and the upper bound for Algorithm 4 (we display $\log_2(T)$ where T is the mean stopping time), together with the 95% confidence intervals. On the right: Mean distance (in 1-norm) between upper and lower bounding states at time 0 for Algorithms 3 and 4, together with the 95% confidence intervals.

This result and Lemma 16 (in Section 5.2) give that the projected chain $\psi(X)$ is between zero and the supremum chain $Y^{\sup}$. If service rate is high, we can wait until the supremum chain $Y^{\sup}$ hits zero. The main advantage is that this provides some solution for the possible decoupling of the system: When $Y^{\sup}$ reaches zero, the projected process is also in state zero, and thus the *only* possible state for the original Markov chain X is also zero. Thus the original chain X couples in $\mathcal{N}$. Hence, we can use coupling from the past for $Y^{\sup}$ until we find a time $-t$ such that $Y^{\sup}_{-t} = 0$, and then, from time $-t$ to time 0, simulate the only trajectory of X starting from state zero (with the same events). This algorithm has quadratic complexity in $|C|$, when $\mu > \sum_i \lambda_i$ (due to Lemma 25 in Appendix C).

We have the following bound for the mean coupling time of the original chain X, which is quadratic with respect to $|C|$:

Theorem 23. *Let τ be the coupling time of the original chain X defined on $\mathcal{N}$, and corresponding to the ATO-POS system with joint services. Assume $\mu > \sum_i \lambda_i$. Then we have:*

$$\mathbb{E}[\tau] \leq \frac{\Lambda}{\mu - \sum_i \lambda_i} |C|.$$

Proof. This is a consequence of Lemma 25 in Appendix C and the fact that the *only* possible state for the chain X when $\psi(X) = 0$ is the state 0: hence if the projected process reaches zero, the original system defined on $\mathcal{N}$ also reaches zero, and so couples. □

7 Further remarks and conclusions

The contribution of our paper is twofold:

- We give perfect sampling algorithms for ATO models with individual and joint replenishments.
- We introduce a new method for perfect sampling, based on aggregation and bounding chains. We apply our new method to the ATO-POS case with joint replenishments, for which the known perfect sampling techniques cannot be applied.

ATO systems with joint services have extremely large state space - its only dimension is $2^I - 1$ - and up to our best knowledge, there is no known efficient solution technique in the literature

(in particular for the POS case that does not have a product form solution). Thus our new perfect sampling method can be of great interest to evaluate their performance, as well as in the optimization algorithms for capacity dimensioning. In most applications, the lost probability is demanded to be very low, thus the conditions in Section 6.7 seem to be natural, under which we can obtain the exact samples of the stationary distribution.

The results in Section 6 can be extended to the TOS case. The arrivals for the TOS case are no longer monotone, so the approach used in Section 6.4 to compute the aggregated envelope chain for the infimum chain has to be used also for the supremum chain. Note that the services are the same as for the POS model.

Up to our knowledge, this is the first time that the aggregation of Markov chains is combined with perfect sampling technique to avoid state space explosion problems. This direction sounds promising for various applications.

A Hitting time to zero for the ATO-POS system with individual state-dependent services (Section 3.2).

We consider a slightly more general model than the one in Section 3: here we allow the service rate μ_i to depend on the whole state x (not only on x_i), provided that services are monotone. This more general setting will be useful to prove Lemma 25 (in Appendix C).

More precisely, for each $1 \leq i \leq I$, let $\mu_i : \mathcal{X} \to [0,1]$, and set $\nu_i := \max_{x\in\mathcal{X}} \mu_i(x)$. Set $\Lambda := \lambda + \sum_i \nu_i$ for the uniformization constant. Arrival events are as before. The set Σ of service events satisfies, for all $x, y \in \mathcal{X}$ and $s \in \Sigma$:

$$x \leq y \;\Rightarrow\; f_P(x,s) \leq f_P(y,s)\,.$$

Let τ_0^P be the mean hitting time that this ATO-POS system reaches the state 0 in $\mathcal{X}$. Previous assumption is satisfied by the model defined in 3.1, so the following proposition also holds for the initial model.

Proposition 24. *Let $\delta := \min_{x\neq 0} \sum_i \mu_i(x) - \sum_i \lambda_i$ and assume $\delta > 0$. Then we have: $\mathbb{E}[\tau_0^P] \leq \frac{\Lambda}{\delta} \sum_{i=1}^I C_i$.*

Proof. Let P be the transition matrix of the ATO-POS model described above. To prove this bound, we will use Lemma 9, with $S = \mathcal{X}$, $U = \{0\}$ and, for all $z \in \mathcal{X}$, $h(z) := \sum_{i=1}^I z_i$. Then we have, for all $y \in \mathcal{X} \setminus \{0\}$:

$$\begin{aligned}
\sum_{z\in\mathcal{X}} P(y,z)h(z) &= \sum_A \frac{\lambda_A}{\Lambda} h\left(y + \sum_{i\in A} \mathbb{1}\{y_i < C_i\}\, e_i\right) + \sum_i \frac{\mu_i(y)}{\Lambda} h(y - e_i) + \sum_i \frac{\nu_i - \mu_i(y)}{\Lambda} h(y) \\
&= \sum_A \frac{\lambda_A}{\Lambda} \left(\sum_i y_i + \sum_{i\in A} \mathbb{1}\{y_i < C_i\}\right) + \sum_i \frac{\mu_i(y)}{\Lambda} \left(\sum_j y_j - 1\right) + \sum_i \frac{\nu_i - \mu_i(y)}{\Lambda} \sum_j y_j \\
&= \left(\sum_A \frac{\lambda_A}{\Lambda} + \sum_i \frac{\nu_i}{\Lambda}\right) h(y) + \sum_A \sum_{i\in A} \mathbb{1}\{y_i < C_i\} \frac{\lambda_A}{\Lambda} - \sum_i \frac{\mu_i(y)}{\Lambda} \\
&\leq h(y) + \sum_A |A| \frac{\lambda_A}{\Lambda} - \sum_i \frac{\mu_i(y)}{\Lambda} = h(y) + \sum_i \frac{\lambda_i}{\Lambda} - \sum_i \frac{\mu_i(y)}{\Lambda} \leq h(y) - \frac{\delta}{\Lambda}.
\end{aligned}$$

Hence the condition (6) of Lemma 9 is proved, and we can apply (7) with $y = (C_1, ..., C_I)$. It follows that $\mathbb{E}[\tau_0] = \mathbb{E}_y[\tau_{\{0\}}] \leq \frac{\Lambda h(y)}{\delta} = \frac{\Lambda}{\delta} \sum_{i=1}^I C_i$ since the time for the system to hit zero is equal to the time for state $y = (C_1, ..., C_I)$ to hit zero, due to monotonicity. □

B Proof of Proposition 12 (Non-monotonicity of the ATO-TOS system with individual replenishments, Section 4.2)

Let $C = (C_1, ..., C_I)$. Let $\preceq$ be a partial order such that f_T is monotone for $\preceq$ (*i.e.* all the events of $\mathcal{E}$ are monotone). We proceed in three steps:

Step 1. Let $x = (x_1, ..., x_I) \in \mathcal{X}$, $x \neq C$. Then we cannot have $x \preceq C$.

Step 2. C is not comparable to any other state.

Step 3. Let $x \in \mathcal{X}$. Then x is not comparable to any other state.

Proof of step 1. Assume by contradiction that $x \preceq C$. Since $x \neq C$, there exists $1 \leq i \leq I$ such that $x_i < C_i$. Without loss of generality, we can assume that $i = 1$. We first show that $v_1 = (C_1 - 1, C_2, ..., C_I)$ also verifies $v_1 \preceq C$. Since $x \preceq C$ and f_T is monotone for $\preceq$, it is enough to find a finite sequence of events that moves x to v_1, and that let C unchanged: For instance, we apply $C_1 - 1 - x_1$ times event $d_{\{1\}}$ and then, for each $2 \leq j \leq I$, we apply $C_j - x_j$ times event $d_{\{j\}}$. Hence $v_1 = (C_1 - 1, C_2, ..., C_I) \preceq C$. Yet the combination $s_2^{(1)}$ followed by $d_{\{1,2\}}$ moves v_1 to C, and C to $v_2 = (C_1, C_2 - 1, ..., C_I)$. Hence, using again the monotonicity of f_T, this leads to: $C \preceq v_2 = (C_1, C_2 - 1, ..., C_I)$. With the same trick ($s_3^{(1)}$ followed by $d_{\{1,3\}}$), we have that $C \preceq v_3 = (C_1, C_2, C_3 - 1, C_4, ..., C_I)$. Yet, starting from $C \preceq v_2 = (C_1, C_2 - 1, ..., C_I)$ and using $s_3^{(1)}$ followed by $d_{\{2,3\}}$ leads to $v_3 \preceq C$. Hence, $v_3 = C$, which is a contradiction.

Proof of step 2. Let $x \neq C$. The same argument as for Step 1 also works, starting from $C \preceq x$ instead of $x \preceq C$. Hence we cannot have $x \preceq C$ nor $C \preceq x$, which proves that C is not comparable to any other state.

Proof of step 3. Let $y = (y_1, ..., y_I) \neq x$: there exists $1 \leq i \leq I$ such that $x_i \neq y_i$. We assume by contradiction that either $y \preceq x$ or $x \preceq y$. Without loss of generality, we can assume that $y \preceq x$. Then $x_i > y_i$. We use the following sequence of events: For each $1 \leq j \leq I$, we apply $C_j - x_j$ times the event $d_{\{j\}}$. This sequence moves x to C and y to y', where y' is defined by: $y'_j = C_j - (x_j - y_j)$ for all $1 \leq j \leq I$. In particular, $y'_i < C_i$, so $y' \neq C$. In addition, the fact that $y \preceq x$ and the monotonicity of f_T for $\preceq$ imply that $y' \preceq C$, which is a contradiction.

C Hitting time to zero for the supremum chain (ATO-POS system with joint returns, Section 6.3)

The next lemma gives the mean hitting time to zero for the supremum chain, using the results of Appendix A and Section 6.3. It is used in the proofs of Theorem 22 (Section 6.6) and Theorem 23 (Section 6.7).

Lemma 25. *Let $\bar{\tau}$ be the time that the supremum chain, starting from $C = (C_1, ..., C_I)$, reaches the state 0 in $\mathcal{X}$. Assume $\mu > \sum_i \lambda_i$. Then we have: $\mathbb{E}[\bar{\tau}] \leq \frac{\Lambda}{\mu - \sum_i \lambda_i} |C|$.*

Proof. We apply Proposition 24 (Appendix A). Lemma 18 and Proposition 19 show that the supremum chain is a particular case of the model presented in Appendix A (with $\mu_i(x) = \mu \hat{x}_i$, $1 \leq i \leq I$). The fact that $\delta = \min_{x \neq 0} \sum_i \mu_i(x) - \sum_i \lambda_i$ is positive comes from the fact that $\min_{x \neq 0} \sum_i \mu_i(x) = \mu \cdot \min_{x \neq 0} \sum_i \hat{x}_i = \mu$ and the hypothesis $\mu > \sum_i \lambda_i$. □

D Proof of Theorem 22 (Stopping time of Algorithm 4, Section 6.6)

The proof of Theorem 22 is based on Lemma 25 (Appendix C) and the following lemma, that gives a bound on the mean hitting time of C_p for the p-th component of the infimum:

Lemma 26. *Let $p \in \mathcal{I}$, and assume $\delta_p := \lambda_p - \mu\left(\sum_{i=1}^{p} C_i - 1\right)$ is positive. Let $\tau^{(p)}$ be the time for the p-th component of $\underline{H}$ to hit C_p (starting from 0). Then: $\mathbb{E}\left[\tau^{(p)}\right] \leq \frac{\Lambda}{\delta_p} C_p$.*

Proof of Lemma 26. In order to give a bound on $\tau^{(p)}$, we have to face two problems: $\underline{H}$ is not the transition function of a Markov chain (it depends on $\overline{H}$), and its projection on the p-th component depends on its whole state (see Lemma 21). That is why we introduce a new Markov chain, defined on $\mathcal{X}$, whose transition function $h^{\inf}$ is a lower bound for $\underline{H}$, and whose projection on the p-th component is also a Markov chain. Indeed, for $m \in \mathcal{X}$, $i \in \mathcal{I}$, $j \in \{1, ..., C_i\}$, we define, for all $p \in \mathcal{I}$ (setting $m' := h^{\inf}(m, r_j^i)$):

- If $p < i$, then $m'_p = m_p$;
- If $p = i$, then $m'_p = m_i - \mathbb{1}\{j \leq m_i\}$;
- If $p > i$, then $m'_p = m_p - \mathbb{1}\{m_p > 0\}$.

In addition, we set $h^{\inf}(m, d_A) = g^{\inf}(m, d_A)$ for all $m \in \mathcal{X}$, $A \subset \mathcal{I}$, $A \neq \emptyset$. Using Lemma 21, we have that, for all $m, M \in \mathcal{X}$, $a \in \mathcal{E}$, $h^{\inf}(m, a) \leq \underline{H}([m, M], a)$. Let $\tau_{\inf}^{(p)}$ be the time for $h^{\inf}$ to hit C_p (starting from 0), then we have that $\tau^{(p)} \leq \tau_{\inf}^{(p)}$. Moreover, we can bound the mean of $\tau_{\inf}^{(p)}$ by:

$$\mathbb{E}\left[\tau_{\inf}^{(p)}\right] \leq \frac{\Lambda}{\delta_p} C_p.$$

The arguments to show this are those of Lemma 10 (3.2) or Proposition 24 (Appendix A): we apply Theorem 9 (3.2) with different parameters, noting that $h^{\inf}$ is also monotone. □

Proof of Theorem 22. Due to condition (iii), the projection of H on $\mathcal{I}_0$ is a Markov chain. Condition (i) allows to apply Lemma 25 to this Markov chain. Condition (ii) allows to apply Lemma 26, and the fact that $\mathcal{I} = \mathcal{I}_0 \cup \mathcal{I}_C$ concludes the proof. □

Acknowledgements

This research is supported by the French National Research Agency grant ANR-12-MONU-0019.

References

[Anselmi and Gaujal, 2011] Anselmi, J. and Gaujal, B. (2011). On the efficiency of perfect simulation in monotone queueing networks. In *IFIP Performance: 29th International Symposium on Computer Performance, Modeling, Measurements and Evaluation*, Amsterdam. ACM Performance Evaluation Review.

[Asmussen and Glynn, 2007] Asmussen, S. and Glynn, P. W. (2007). *Stochastic simulation: algorithms and analysis*, volume 57 of *Stochastic Modelling and Applied Probability*. Springer, New York.

[Bušić et al., 2012a] Bušić, A., Gaujal, B., and Pin, F. (2012a). Perfect sampling of Markov chains with piecewise homogeneous events. *Performance Evaluation*, 69(6):247–266.

[Bušić et al., 2008] Bušić, A., Gaujal, B., and Vincent, J.-M. (2008). Perfect simulation and non-monotone markovian systems. In *Valuetools'08*, Athens, Grece.

[Bušić et al., 2012b] Bušić, A., Vliegen, I., and Scheller-Wolf, A. (2012b). Comparing Markov chains: Aggregation and precedence relations applied to sets of states, with applications to assemble-to-order systems. *Mathematics of Operations Research*, 37(2):259–287.

[Dayanik et al., 2003] Dayanik, S., Song, J.-S., and Xu, S. H. (2003). The effectiveness of several performance bounds for capacitated production, partial-order-service, assemble-to-order systems. *Manufacturing & Service Operations Management*, 5(3):230–251.

[Dopper et al., 2006] Dopper, J., Gaujal, B., and Vincent, J.-M. (2006). Bounds for the coupling time in queueing networks perfect simulation. In *Numerical Solutions for Markov Chains (NSMC'06)*, pages 117–136, Charleston. The 2006 A.A. Markov Anniversary Meeting (MAM 2006).

[Huber, 2004] Huber, M. (2004). Perfect sampling using bounding chains. *Ann. Appl. Probab.*, 14(2):734–753.

[Kelly, 1991] Kelly, F. P. (1991). Loss networks. *The Annals of Applied Probability*, 1(3):319–378.

[Kendall and Møller, 2000] Kendall, W. S. and Møller, J. (2000). Perfect simulation using dominating processes on ordered spaces, with application to locally stable point processes. *Advances in Applied Probability*, 32(3):844–865.

[Levin et al., 2009] Levin, D. A., Peres, Y., and Wilmer, E. L. (2009). *Markov chains and mixing times*. American Mathematical Society, Providence, RI. With a chapter by J. G. Propp and D. B. Wilson.

[Li and Xu, 2000] Li, H. and Xu, S. (2000). On the dependence structure and bounds of correlated parallel queues and their applications to synchronized stochastic systems. *Journal of Applied Probability*, 37(4):1020–1043.

[Louth et al., 1994] Louth, G., Mitzenmacher, M., and Kelly, F. (1994). Computational complexity of loss networks. *Theoretical Computer Science*, 125(1):45–59.

[Propp and Wilson, 1996] Propp, J. G. and Wilson, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms*, 9(1-2):223–252.

[Song et al., 1999] Song, J.-S., Xu, S. H., and Liu, B. (1999). Order-fulfillment performance measures in an assemble-to-order system with stochastic leadtimes. *Operations Research*, 47(1):131–149.

[Song and Zipkin, 2003] Song, J.-S. and Zipkin, P. (2003). Supply chain operations: Assemble-to-order systems. In de Kok, A. and Graves, S., editors, *Supply Chain Management: Design, Coordination and Operation*, volume 11 of *Handbooks in Operations Research and Management Science*, chapter 11, pages 561–596. North-Holland.

[Vliegen and van Houtum, 2009] Vliegen, I. M. H. and van Houtum, G. J. (2009). Approximate evaluation of order fill rates for an inventory system of service tools. *International Journal of Production Economics*, 118(1):339–351.

[Wilson, 1998] Wilson, D. B. (1998). Perfectly random sampling with Markov chains, annotated bibliograpy. *DIMACS Series in Discret Mathematics and Theoretical Computer Science*, 41. http://dimacs.rutgers.edu/ dbwilson/exact.html/.

[Xu, 1999] Xu, S. (1999). Structural analysis of a queueing system with multiclasses of correlated arrivals and blocking. *Operations Research*, 47(2):264–276.

[Xu, 2002] Xu, S. (2002). Dependence analysis of assemble-to-order systems. In Song, J.-S. and Yao, D., editors, *Supply Chain Structures: Coordination, Information and Optimization*, chapter 11. Kluwer Academic, Boston.